

IMMERSED TURNOVERS IN HYPERBOLIC 3-ORBIFOLDS

SHAWN RAFALSKI

ABSTRACT. We show that any immersion, which is not a covering of an embedded 2-orbifold, of a totally geodesic hyperbolic turnover in a complete orientable hyperbolic 3-orbifold is contained in a hyperbolic 3-suborbifold with totally geodesic boundary, called the “turnover core,” whose volume is bounded from above by a function depending only on the area of the given turnover. Furthermore, we show that, for a given type of turnover, there are only finitely many possibilities for the turnover core. As a corollary, if the volume of a complete orientable hyperbolic 3-orbifold is at least 2π and if the fundamental group of the orbifold contains the fundamental group of a hyperbolic turnover (i.e., a triangle group), then the orbifold contains an embedded hyperbolic turnover.

1. INTRODUCTION

It is well known that the thrice-punctured sphere is the only orientable hyperbolic surface which is rigid, in the sense that it admits a unique complete hyperbolic structure. The analogue of this surface in the orbifold setting is the hyperbolic turnover. A *hyperbolic turnover* is the double, along the boundary, of a hyperbolic triangle whose interior angles are integer submultiples of π . As an orbifold, a hyperbolic turnover is topologically a 2-sphere with three cone points whose orders correspond to the submultiples of π in the associated hyperbolic triangle. Like the thrice-punctured sphere, hyperbolic turnovers admit a unique complete hyperbolic structure. The goal of this paper is to prove the following theorem, which can be viewed either as a finiteness result or as the turnover analogue of several well-known theorems from classical 3-manifold topology (see Corollaries 1.2 and 1.3 below).

Theorem 1.1. *Let $f: \mathcal{T} \rightarrow Q$ be a totally geodesic (equivalently, π_1 -injective) immersion of a compact, hyperbolic turnover $\mathcal{T} = \mathcal{T}(p, q, r)$ in an orientable hyperbolic 3-orbifold Q . Assume that $f(\mathcal{T})$ does not cover an embedded turnover or an embedded triangle with mirrored sides. Then Q contains a finite (possibly empty) collection $\{\mathcal{T}_i\}$ of embedded, pairwise disjoint, totally geodesic hyperbolic turnovers (and*

totally geodesic hyperbolic triangles with mirrored sides) satisfying the following:

- (1) $f(\mathcal{T}) \cap \mathcal{T}_i = \emptyset$ for each i
- (2) *The number of turnovers (and triangles with mirrored sides) in the collection $\{\mathcal{T}_i\}$ is bounded above by a function of p, q and r*
- (3) *If n is the order of a cone point of any turnover in the collection $\{\mathcal{T}_i\}$, then*

$$n \in \{2, 3, \dots, 9, p, q, r, 2p, 2q, 2r\}.$$

- (4) *If π/n is an angle of a triangle with mirrored sides in the collection $\{\mathcal{T}_i\}$, then*

$$n \in \{2, 3, \dots, 9, p, q, r, 2p, 2q, 2r\}.$$

- (5) *If Q' is the component of $Q - \cup_i \mathcal{T}_i$ which contains $f(\mathcal{T})$, then the metric closure of Q' is a small hyperbolic 3-orbifold with (possibly empty) totally geodesic boundary. If $\partial Q'$ is not empty, then $\text{Vol}(Q') < H \cdot \text{Area}(\mathcal{T})$, where $H = 1.199678\dots$ is the positive solution of the equation $x = \coth x$. If $\partial Q'$ is empty, then $\text{Vol}(Q') < \text{Area}(\mathcal{T})$.*

Furthermore, for a given (p, q, r) -turnover, there are only finitely many possibilities for the orbifold Q' described above.

Because the area of any hyperbolic turnover is bounded above by 2π , we have the following

Corollary 1.2. *(The Turnover Theorem) Let Q be an orientable hyperbolic 3-orbifold with $\text{Vol}(Q) \geq 2\pi$. Then Q contains an embedded turnover (or an embedded triangle with mirrored sides), if it contains an immersed turnover.*

Dunbar [11] showed that every compact, irreducible, atoroidal 3-orbifold can be split (uniquely, up to isotopy) along a system of essential, pairwise non-parallel hyperbolic turnovers into pieces which contain no essential (embedded) turnovers. The next corollary, which follows from the arguments of Section 5, says that the Dunbar decomposition is the turnover analogue of the JSJ-decomposition of a 3-manifold.

Corollary 1.3. *Let Q be a compact, irreducible, orientable, atoroidal 3-orbifold. Then any immersion $f: \mathcal{T} \rightarrow Q$ of a hyperbolic turnover into Q is homotopic into a unique component of the Dunbar decomposition, up to parallel boundary components of the decomposition.*

The motivation for Theorem 1.1 begins with a question which Kirby attributes to Martin [17, Problem 3.70]:

Question 1.4. *Given a Kleinian group Γ and a turnover subgroup T with invariant plane Π in \mathbb{H}^3 , is it true that for all $\gamma \in \Gamma$ either $\gamma\Pi = \Pi$ or $\gamma\Pi \cap \Pi = \emptyset$? Equivalently, is it true that the turnover $\mathcal{T} = \Pi/T$ covers an embedded turnover in the orbifold \mathbb{H}^3/Γ ?*

As it turns out, the answer to this question is no, although the author believes that this must have been known to Martin at the time that Kirby added this question to his list of problems in low dimensional topology. Consider the tetrahedron in Figure 1. The




FIGURE 1. A compact hyperbolic tetrahedron which yields a 3-orbifold with an immersed turnover

integers at the edges indicate the dihedral angles as submultiples of π , so, for instance, an edge labeled 2 has a dihedral angle of $\pi/2$. It is known that this tetrahedron can be realized in hyperbolic 3-space, and that the group generated by reflections in its faces is a discrete group of isometries. Let Γ_3 be the orientation-preserving subgroup of index two inside this reflection group. Consider the face ABC . It is not difficult to see that this face is a hyperbolic triangle with the angles $\pi/2, \pi/4$, and $\pi/5$, as indicated in the figure. Consider the centralizer, in the reflection group, of reflection in this face. Baskan and Macbeath [2] proved that the orientation-preserving subgroup of this centralizer is a turnover subgroup of Γ_3 . (A *turnover group* is the orbifold fundamental group of a turnover.) The corresponding geometric turnover is immersed in the quotient 3-orbifold, because

the order 3 edge AD meets the plane containing the face ABC at an oblique angle.

There are nine compact hyperbolic tetrahedra for which the group of isometries generated by edge rotations is discrete. In all nine cases, the corresponding hyperbolic 3-orbifolds contain immersed turnovers. The other examples are not as easily seen as in the Baskan and Macbeath example above, however, for this tetrahedron is the only one with a triangular face whose interior angles are integer submultiples of π . Maclachlan has classified the immersed turnovers in these “tetrahedral” groups for eight of the nine cases [20], although he makes an error in one of the cases which we will rectify (Proposition 9.3). In any event, all of these counterexamples to Question 1.4 are “small” in several respects. We will consider them in some detail in Section 9.4.

On the other hand, Martin has proved the affirmative answer to his question for at least one class of hyperbolic turnover [21]. Specifically, he showed that $(2, 3, p)$ turnovers, where $p \geq 7$, are always embedded in a hyperbolic 3-orbifold. In joint work with Gehring [15], this result was utilized as part of a general program for locating the minimal co-volume Kleinian group, which is conjectured to be the orientation-preserving index two subgroup of a particular Coxeter reflection group (in fact, the reflection group corresponds to one of the nine tetrahedral examples mentioned above).

Knowing that a turnover is embedded in a hyperbolic 3-orbifold can be used to find lower bounds for the volume of the 3-orbifold, by considering embedded δ -neighborhoods of the turnover (e.g. [21, Theorem 1.11]). Furthermore, Kleinian groups which contain a turnover subgroup (or equivalently, hyperbolic 3-orbifolds which contain hyperbolic turnovers) frequently turn up as candidates for the solutions to several types of extremal problems in hyperbolic geometry. For instance, recent work of Gehring and Martin [16] attempts to determine the Margulis constant for \mathbb{H}^3 , and the elements of one class of groups they consider (so called (p, q, r) -Kleinian groups) frequently contain turnover subgroups.

Finally, because turnovers have the unique property of rigidity among orientable 2-orbifolds, they act as shields to the effects of deformation of the hyperbolic structure of the ambient hyperbolic 3-orbifold. Specifically, if a turnover \mathcal{T} separates off a cusp C of the 3-orbifold Q from the rest of Q , then any hyperbolic Dehn surgery performed on C affects only the geometry of the component of $Q - \mathcal{T}$ which

contains C . This is a specific instance of a phenomenon known as *geometric isolation* [26], [7]. It follows from Corollary 1.2 that once the volume of a cusped hyperbolic 3-orbifold Q is large enough, then knowing its fundamental group contains a turnover subgroup is equivalent to knowing that there are pieces of Q which are left alone by many deformations. In light of all of these considerations, a complete classification of immersed turnovers in hyperbolic 3-orbifolds would be quite useful. It is the hope of the author that Theorem 1.1 is a first step toward that classification.

Acknowledgments. *The author is grateful to his advisor Ian Agol for giving so freely of his time to discuss mathematics in all its forms. Special thanks also go to Marc Culler and Peter Shalen, who generously provided their time and suggestions for the work contained herein. Finally, the author thanks the referee for invaluable feedback.*

2. DEFINITIONS AND NOTATION

An n -orbifold O is a metrizable topological space in which every point has a neighborhood which is diffeomorphic either to the quotient of \mathbb{R}^n by a finite group action or to the quotient of $\mathbb{R}^{n-1} \times [0, \infty)$ by a finite group action. Points with neighborhoods modeled on the latter type of quotients make up the *boundary* ∂O of the n -orbifold, which is itself an $(n-1)$ -orbifold. An orbifold is called *geometric* if its interior is diffeomorphic to the quotient of a model geometric space by a discrete group of isometries. We will not define the term *model geometric space*, but will rather point out that the n -dimensional sphere, Euclidean space, and hyperbolic space (denoted, respectively, by \mathbb{S}^n , \mathbb{E}^n , and \mathbb{H}^n) are the examples of model geometric spaces in which we will be interested. There are several extensive references for and introductions to the definitions in this section [4], [10].

In all of what follows, Q is a complete orientable hyperbolic 3-orbifold. That is, Q is the quotient \mathbb{H}^3/Γ of hyperbolic 3-dimensional space by a Kleinian group Γ . A *Kleinian group* is a discrete and non-elementary subgroup of $\mathrm{PSL}_2(\mathbb{C})$, where we identify $\mathrm{PSL}_2(\mathbb{C})$ with the group of orientation-preserving isometries of \mathbb{H}^3 via the Poincaré extension. Alternatively, Γ may be thought of as a subgroup of $\mathrm{PSL}_2(\mathbb{C})$ whose action on \mathbb{H}^3 is properly discontinuous (although not necessarily free). The action of a group G on a space X is *properly discontinuous* if any compact subset $K \subset X$ is taken completely off of itself by all but only finitely many elements of G . We will refer to Γ as the (*orbifold*)

fundamental group of Q , and sometimes denote it by $\pi_1(Q)$. Denote the covering projection by $\Phi: \mathbb{H}^3 \rightarrow Q$.

Let $2 \leq p \leq q \leq r$ be positive integers satisfying $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} < 1$ (in generalizations of the terminology presented here it is allowed for p, q , and r to take on the value ∞). Then there is a hyperbolic triangle $\Delta = \Delta(p, q, r)$, unique up to isometry, with interior angles $\frac{\pi}{p}$, $\frac{\pi}{q}$ and $\frac{\pi}{r}$. The group generated by reflections in the sides of Δ is a discrete subgroup of isometries of 2-dimensional hyperbolic space \mathbb{H}^2 . Let $T(p, q, r)$ be the unique normal subgroup of index two inside this reflection group which acts on \mathbb{H}^2 by orientation-preserving isometries. Then $T(p, q, r)$ is generated by rotations by $\frac{2\pi}{p}$, $\frac{2\pi}{q}$ and $\frac{2\pi}{r}$ around the corresponding vertices of Δ (in fact, any pair of these rotations generates this group). Call this a *turnover group* of $\text{Isom}(\mathbb{H}^2)$. (A remark: These are commonly referred to as *triangle groups*, and while the author has a great deal of respect for tradition, he feels somewhat compelled to refer to these groups by the likeness of their associated 2-dimensional geometric objects.) A *turnover subgroup* of a Kleinian group Γ is a subgroup which is isomorphic to some $T(p, q, r)$. We will denote it by $T = T(p, q, r)$.

A turnover subgroup $T(p, q, r)$ of a Kleinian group Γ is generated by three elliptic elements γ_p , γ_q and γ_r of orders p , q and r , respectively (or any pair of these, see above). By the rigidity of turnover groups in $\text{Isom}(\mathbb{H}^3)$ [22, Chapter IX.C], there is a geodesic plane $\Pi_T \subset \mathbb{H}^3$ which is invariant under the action of T , and on which T acts as a standard turnover group, with fundamental domain consisting of two $\Delta(p, q, r)$ triangles. Thus, by the Gauss-Bonnet theorem, the area of $\mathcal{T} := \Pi_T/T$ is $2\pi(1 - (\frac{1}{p} + \frac{1}{q} + \frac{1}{r}))$, and \mathcal{T} is a 2-dimensional space, homeomorphic to a 2-sphere, with a Riemannian metric of constant curvature -1 in the complement of three points and such that each of these three points has a neighborhood isometric to a hyperbolic cone. We will refer to such a \mathcal{T} as a *hyperbolic turnover*, and denote by $f: \mathcal{T} \rightarrow Q$ the restriction to \mathcal{T} of the covering map $\mathbb{H}^3/T \rightarrow Q$, and we will call $f(\mathcal{T}) \subset Q$ an *immersion* of a hyperbolic turnover in the hyperbolic 3-orbifold Q . Notice that $\Phi^{-1}(f(\mathcal{T})) = \bigcup_{\gamma \in \Gamma} \gamma \Pi_T$ is a union of geodesic planes in \mathbb{H}^3 . We will say in this case that the immersion is *totally geodesic*. Thus, an immersion of a hyperbolic turnover $f: \mathcal{T} \rightarrow Q$ is totally geodesic if and only if the map is injective on the level of fundamental groups. We will call the non-orientable 2-orbifold doubly

covered by a turnover a *triangle with mirrored sides*, or often just a *triangle* if the context is clear.

There are turnovers with spherical and Euclidean structures as well, which are obtained as index two orientation-preserving subgroups of discrete groups generated by reflections in the sides of spherical or Euclidean triangles, respectively. Using the notation of triples to denote the submultiples of π in the given triangle, there are four types of spherical turnover: $(2, 3, 3)$, $(2, 3, 4)$, $(2, 3, 5)$, and $(2, 2, n)$ for $n \geq 2$. The first three correspond to the quotient of the 2-sphere \mathbb{S}^2 by the orientation-preserving isometries of the regular tetrahedron, cube, and dodecahedron, respectively. The last is the quotient of \mathbb{S}^2 by the dihedral group generated by two order 2 rotations whose axes meet in an angle of π/n . If we add to this list both \mathbb{S}^2 and the quotient of \mathbb{S}^2 by a cyclic group of order n , then it is well known that every point in a complete orientable hyperbolic 3-orbifold has a neighborhood which is isometric to the cone on one of these six types of spherical 2-orbifolds.

The Euclidean turnovers are $(2, 3, 6)$, $(2, 4, 4)$, $(3, 3, 3)$, and $(2, 2, \infty)$. The first three correspond to the doubles of the standard Euclidean triangles, and the last is an open-ended noncompact “pillowcase” (i.e., the double of a finite-width infinite half-strip in the plane). The other Euclidean 2-orbifolds we will consider are the torus and the quotient “pillow” obtained from the torus by an involution with four fixed points.

Let \mathcal{O} be a compact 3-orbifold, possibly with boundary. We use the term *Haken ball* to refer to the quotient of a compact 3-ball by a finite group of isometries. We say \mathcal{O} is *irreducible* if every embedded spherical 2-suborbifold of \mathcal{O} bounds a Haken ball in \mathcal{O} , and *atoroidal* if every π_1 -injective map of a torus, pillow, or Euclidean turnover into \mathcal{O} is parallel into a boundary component (parallel in the sense that the 2-orbifold cuts off a product neighborhood of a boundary component). An arbitrary 3-orbifold is called *atoroidal* if it is diffeomorphic to the interior of an atoroidal 3-orbifold with boundary. A 2-orbifold F in \mathcal{O} is *compressible* if either F is a spherical 2-orbifold which bounds a Haken ball or if there is a homotopically nontrivial curve in F which bounds a disk quotient (i.e., an *orbifold disk*) in \mathcal{O} . We say F is *incompressible* otherwise. Similarly, there is a relative notion of ∂ -*compressibility* and ∂ -*incompressibility* (whose exact definition we will not require). We call F *essential* if it is incompressible, ∂ -incompressible, and not parallel into a boundary component. We

call a compact irreducible 3-orbifold *Haken* if it is either a Haken ball, or a turnover crossed with an interval, or if it contains an essential 2-suborbifold but contains no essential turnover. A compact irreducible 3-orbifold is called *small* if contains no essential 2-suborbifolds and has (possibly empty) boundary consisting only of turnovers. The definitions of *essential*, *Haken*, and *small* extend to arbitrary 3-orbifolds in the same way as the definition of atoroidal. Finally, we say a n -orbifold is *good* if it is covered by a n -manifold. All orbifolds considered in the paper here are assumed to be good.

3. ANALYSIS OF THE COMPLEMENT OF THE TURNOVER

The idea behind the proof of Theorem 1.1 is essentially due to Cooper [9], that is, to cut the 3-orbifold along the turnover and then bound the volumes of the complementary pieces. In this way, we will bound from above the volume of the 3-orbifold.

3.1. Preliminaries. To begin, let $T \leq \Gamma \cong \pi_1(Q)$ be a turnover subgroup of the fundamental group of the complete orientable hyperbolic 3-orbifold Q , and let $f: \mathcal{T} \rightarrow Q$ be the associated isometric immersion. Let $\Phi: \mathbb{H}^3 \rightarrow Q$ be the covering projection. We will require the following well known result. Its proof follows from the proper discontinuity of the action of Γ and the Margulis lemma.

Proposition 3.1. *Let $f: S \rightarrow Q$ be a totally geodesic immersion of any 2-orbifold S of finite area in a hyperbolic 3-orbifold Q . Let $\Gamma < \text{Isom}(\mathbb{H}^3)$ be the fundamental group of Q , and let $\Gamma_S \leq \Gamma$ be an isomorphic copy of $\pi_1(S)$ stabilizing a geodesic plane $\Pi \subset \mathbb{H}^3$. Then the collection of planes $\Phi^{-1}(f(S)) = \bigcup_{\gamma \in \Gamma} \gamma \Pi$ is locally finite, in the sense that any compact set $K \subset \mathbb{H}^3$ meets only finitely many planes in the collection.*

Let $\bigcup_{j \in \mathcal{J}} S_j$ be the disjoint union of the components of $Q - f(\mathcal{T})$, and let Σ_j be a connected component of $\Phi^{-1}(S_j)$. Define Γ_j to be the stabilizer of Σ_j in Γ , and let $\Lambda(\Gamma_j)$ denote the limit set of Γ_j , that is, the collection of accumulation points of the orbit space $\Gamma_j x$ on the sphere at infinity S_∞^2 of \mathbb{H}^3 , where $x \in \mathbb{H}^3$ is any point. Note that $\Gamma_j \cong \pi_1(S_j)$. Denote the set $\mathbb{H}^3 \cup S_\infty^2$ by $\overline{\mathbb{H}}^3$. The *convex hull* of $\Lambda(\Gamma_j)$, which is denoted by $CH(\Lambda(\Gamma_j))$, is defined to be the smallest convex set in $\overline{\mathbb{H}}^3$ containing $\Lambda(\Gamma_j)$ (equivalently, $CH(\Lambda(\Gamma_j))$ is the intersection of all the closed half-spaces of $\overline{\mathbb{H}}^3$ which contain $\Lambda(\Gamma_j)$).

The set $CC(S_j) := \Phi(CH(\Lambda(\Gamma_j)) \cap \mathbb{H}^3)$ is called the *convex core* of S_j . It is the smallest convex subset of S_j which carries its fundamental group.

Let $C \subset \mathbb{H}^3$ be a closed convex set which is invariant by a Kleinian group G , and for this paragraph only denote by \overline{C} the closure of C in $\overline{\mathbb{H}^3}$. Then there is well-defined, G -invariant projection $r: \overline{\mathbb{H}^3} - \overline{C} \rightarrow C$ which is given by nearest point retraction. This is defined as follows. If $x \in \overline{\mathbb{H}^3} - \overline{C}$, then take $B(x, \rho)$ to be a ball of radius ρ (or horoball, if $x \in S_\infty^2$) around x and expand this ball until it meets C . This unique first point of intersection is $r(x)$. It is the intersection of a ball (or horoball) around x with a supporting half-plane for the convex set C . The inverse image of the frontier of C , $r^{-1}(\text{fr}(C)) = r^{-1}(C) \cap \overline{\mathbb{H}^3} - \overline{C}$, is just the union, over all $z \in \text{fr}(C)$, of rays beginning at z and which are perpendicular to a supporting half-plane at z . Furthermore, if C is not contained in a proper hyperbolic subspace, then $\text{fr}(C)$ is homeomorphic to $S_\infty^2 - \overline{C}$. The proofs of these facts are all contained in the work of Epstein and Marden [13, Sections 1.2–1.4].

In particular, when the convex set in question is $CH(\Lambda(\Gamma_j))$, we can use these facts about the nearest point retraction to determine the shape of Σ_j . We will conduct the analysis based on the cardinality of the limit set $\Lambda(\Gamma_j)$, which contains 0, 1, 2 or infinitely many points [22, Chapter II.D]. We call a Kleinian group *elementary* or *nonelementary* depending on whether or not it has a finite limit set. Elementary Kleinian groups are those which contain an abelian subgroup of finite index.

3.2. Case: $|\Lambda(\Gamma_j)| = 0$. According to the classification of elementary Kleinian groups [22, Chapter V], Γ_j is a finite group isomorphic to the trivial group, a cyclic group, a dihedral group, or the group of orientation-preserving rigid motions of one of the five Platonic solids. In particular, Γ_j fixes a point in $z \in \Sigma_j$, and Σ_j is the union of all geodesic rays from z that miss $\Phi^{-1}(f(\mathcal{T}))$ and all segments from z to $\Phi^{-1}(f(\mathcal{T}))$ whose interiors miss $\Phi^{-1}(f(\mathcal{T}))$.

The set $\Phi^{-1}(f(\mathcal{T}))$ is locally finite and its complement in \mathbb{H}^3 is homeomorphic to a collection of open balls (the latter fact is most easily seen using the projective ball model of \mathbb{H}^3). Thus, $\overline{\Sigma_j}$ is a solid hyperbolic polyhedron P (henceforth, we will use the bar to denote closure in the ambient space, either \mathbb{H}^3 or Q). Because \mathcal{T} has finite diameter and finite area, the closure $\overline{S_j}$ of a component S_j of $Q - f(\mathcal{T})$ with finite fundamental group must be compact. Therefore, since P

covers \overline{S}_j by the action of the finite group Γ_j , P must be a compact polyhedron. The isoperimetric inequality for hyperbolic 3-space (e.g. [6, Section 9], [8, Section 6.4]) implies that the ratio of surface area to volume enclosed is always greater than 2. Applying this to Σ_j gives $\text{Vol}(\Sigma_j) < \text{Area}(\partial\overline{\Sigma}_j)/2$, and this inequality passes to the quotient:

Lemma 3.2. *If S_j is a component of the complement of the immersed turnover $f(\mathcal{T})$ in Q and $\pi_1(S_j)$ is elementary with empty limit set, then $CC(S_j)$ is a point and*

$$\text{Vol}(S_j) < \frac{\text{Area}(\partial\overline{S}_j)}{2}.$$

3.3. Case: $|\Lambda(\Gamma_j)| = 1$. In this case, Γ_j fixes a point on the sphere at infinity S_∞^2 , which we will assume to be the point ∞ in the upper half-space model of \mathbb{H}^3 , and the connected component $\overline{\Sigma}_j$ is the union of all geodesics segments from ∞ to $\Phi^{-1}(f(\mathcal{T}))$ which are contained in $\overline{\Sigma}_j$. It remains to determine what this collection looks like. Again we appeal to the classification of elementary Kleinian groups. In particular, the maximal parabolic subgroup of Γ_j has rank 1 or rank 2.

We begin with the following observation. In both the rank 1 and rank 2 cases, the classification of elementary Kleinian groups provides that Γ_j stabilizes any horoball centered at ∞ . We would like to see that $\overline{\Sigma}_j$ must project “vertically” onto such a horoball. If not, then there must be vertical planes in $\Phi^{-1}(f(\mathcal{T}))$ (that is, planes containing the point ∞) which cut out some subset of a horoball centered at ∞ in such a way that the subset is invariant by Γ_j . For example, Σ_j might have two vertical “walls” which cut out a “slab” on which Γ_j acts by translations. See Figure 2, which is meant to illustrate this possibility. Each of the polygonal pieces should be imagined as part of hemisphere that is perpendicular to the bounding plane (i.e., a polygon in a geodesic plane of \mathbb{H}^3). We observe, however, that this possibility cannot occur for the simple reason (observed above in the previous case) that a compact turnover has finite diameter. This observation, combined with the fact that Σ_j is connected and stabilized by Γ_j , proves the following:

Lemma 3.3. *If S_j is a component of $Q - f(\mathcal{T})$ and $\pi_1(S_j)$ is elementary with limit set a single point, then the closure of a connected component Σ_j of $\Phi^{-1}(S_j)$ is isotopic to a horoball centered at the fixed point of $\Gamma_j = \text{Stab}_\Gamma(\Sigma_j)$, with the isotopy given by projection toward this fixed point.*




FIGURE 2. A vertical “slab” cut out of a horoball by $\Phi^{-1}(f(\mathcal{T}))$, with direction of the action of Γ_j indicated by the arrow in the bounding plane

If the rank of the maximal parabolic subgroup of Γ_j is 1, then Γ_j is isomorphic to either \mathbb{Z} or a Euclidean $(2, 2, \infty)$ turnover group, each of these groups acting on horoballs centered at ∞ . In either case, we will derive a contradiction. Observe that a fundamental domain for the action of Γ_j on a horosphere centered at ∞ has infinite area. At this point, we have not shown that the collection of planes $\Phi^{-1}(f(\mathcal{T}))$ is connected. Consequently, it could happen that $\partial\overline{\Sigma}_j$ does not project onto the whole horosphere. If the projection *is* onto the whole horosphere, then the isotopy of Lemma 3.3 maps a fundamental domain for Γ_j in $\partial\overline{\Sigma}_j$ onto a set of infinite area in a fundamental domain for Γ_j in a horosphere centered at ∞ . This gives an immediate contradiction, because $f(\mathcal{T})$ has finite area in Q . But if the mapping is onto a set of *finite* area in a fundamental domain for Γ_j in a horosphere centered at ∞ , then we still have a contradiction, as follows.

The subgroup $\Gamma_j < \Gamma$ acts on horoballs so that the quotient component $S_j = \Sigma_j/\Gamma_j \subset Q - f(\mathcal{T})$ is an open convex solid torus or solid pillow. The boundary $\partial\overline{S}_j$, which is composed of pieces of $f(\mathcal{T})$, is a compact set in Q with no boundary. As a consequence, it can have no

ends. We must therefore have that $\partial\overline{S}_j$ lifts to a set in a fundamental domain for the action of Γ_j which has no ends. Now we have that $\Phi^{-1}(f(\mathcal{T}))$ is a union of geodesic planes which cuts out a collection of open 3-balls in \mathbb{H}^3 , and we also have that a fundamental domain for Γ_j in a horosphere centered at ∞ has infinite area. These two facts imply that, if the image (under the above isotopy) of a fundamental domain for Γ_j in $\partial\overline{S}_j$ has finite area in a fundamental domain in a horosphere, then $\partial\overline{S}_j$ must lift to a set with ends in a fundamental domain for Γ_j . But this cannot happen. We conclude that the rank 1 case cannot occur.

If the rank of the maximal parabolic subgroup of Γ_j is 2, then by the classification of elementary Kleinian groups, Γ_j is isomorphic to one of the compact Euclidean turnover groups $((2, 3, 6), (2, 4, 4), \text{ or } (3, 3, 3))$, or the fundamental group of a pillowcase or torus. When we project to Q , we see that $f(\mathcal{T})$ cuts out a neighborhood of a finite-volume cusp, which is called *rigid* in the Euclidean turnover case and *non-rigid* otherwise. See Figure 3 for a possible illustration in the torus case.




FIGURE 3. The fundamental domain for the action of Γ_j in the rank 2 case

We will eventually see that such a cusp must be rigid (this is, in essence, due to the fact that a hyperbolic turnover has a unique hyperbolic structure). In this case, that is, the compact Euclidean turnover group case, we can bound the volume of a fundamental domain in \overline{S}_j for Γ_j in terms of the area of (a portion of) its boundary. We let F

denote such a fundamental domain, and we choose F to be a union of totally geodesic simplices. The set $\partial F \cap \partial \bar{\Sigma}_j$ consists of totally geodesic simplices, each point of which has a unique geodesic segment connecting it to ∞ (that is, an observer at ∞ can see all of $\partial F \cap \partial \bar{\Sigma}_j$). Observe that $\text{Area}(\partial F \cap \partial \bar{\Sigma}_j) = \text{Area}(\partial \bar{S}_j)$. The following result is presumably well known. We provide the argument for completeness.

Lemma 3.4. $\text{Vol}(F) < \frac{1}{2} \text{Area}(\partial F \cap \partial \bar{\Sigma}_j)$.

Proof of 3.4. Since $\partial F \cap \partial \bar{\Sigma}_j$ is made up of totally geodesic triangles, it suffices to prove the result when F is just the region above a geodesic triangle in the upper half-space model of \mathbb{H}^3 . See Figure 4. After applying an isometry of \mathbb{H}^3 , we may assume that the hemisphere representing the geodesic plane containing our triangle has Euclidean radius 1. Let Δ be the triangle in the bounding plane which is the Euclidean orthogonal projection of $\partial F \cap \partial \bar{\Sigma}_j$. Note that the equatorial disk of our unit hemisphere represents a copy of the projective model for \mathbb{H}^2 , and Δ is an isometric copy of $\partial F \cap \partial \bar{\Sigma}_j$ (in the projective model) under the projection. We have the following calculation

$$\begin{aligned} \text{Vol}(F) &= \iint_{(x,y) \in \Delta} \left(\int_{\sqrt{1-(x^2+y^2)}}^{\infty} \frac{dz}{z^3} \right) dx dy \\ &= \frac{1}{2} \iint_{(x,y) \in \Delta} \frac{dx dy}{1 - (x^2 + y^2)} < \frac{1}{2} \iint_{(x,y) \in \Delta} \frac{dx dy}{(1 - (x^2 + y^2))^{\frac{3}{2}}} = \frac{1}{2} \text{Area}(\partial F \cap \partial \bar{\Sigma}_j). \end{aligned}$$

The final equality is obtained by recognizing the integrand of the last double integral as the area element for the projective model of \mathbb{H}^2 . This proves the lemma. 3.4

Observe that this result applies to any finite-volume cusp neighborhood which is cut out from Q by $f(\mathcal{T})$. By Corollary 5.5, however, the non-rigid cusp neighborhoods will not occur. In particular, we have the following:

Lemma 3.5. *If S_j is a component of $Q - f(\mathcal{T})$ and $\pi_1(S_j)$ is elementary with 1 limit point, then S_j is a finite-volume, rigid cusp neighborhood in Q , and*

$$\text{Vol}(S_j) < \frac{\text{Area}(\partial \bar{S}_j)}{2}.$$




FIGURE 4. The region, in upper half-space, above a geodesic triangle

3.4. **Case:** $|\Lambda(\Gamma_j)| = \infty$. Since $\Lambda(\Gamma_j)$ is contained in one “side” of S_∞^2 determined by any plane $\Pi \subset \Phi^{-1}(f(\mathcal{T}))$ which contains a facet of $\overline{\Sigma}_j$, it follows that $CH(\Lambda(\Gamma_j)) \neq \mathbb{H}^3$. So $CH(\Lambda(\Gamma_j))$ projects to $CC(S_j)$ in Q as either a totally geodesic hyperbolic 2-suborbifold or a hyperbolic 3-suborbifold with a hyperbolic 2-orbifold in its boundary which separates $f(\mathcal{T})$ from $CC(S_j)$ [30, Proposition 8.5.1], [13, Theorem 1.12.1]. The flow induced by the gradient of the function which gives the distance from a point to $CH(\Lambda(\Gamma_j))$ determines a product structure (e.g. [1, Appendix A])

$$\overline{S}_j - CC(S_j) \cong \partial \overline{S}_j \times [0, 1).$$

We make some remarks to give an idea of what is going on here. Suppose that $CH(\Lambda(\Gamma_j))$ is a geodesic plane. If Γ_j contains no elliptic element of order two whose axis is contained in this plane, i.e., if Γ_j is an orientation-preserving Fuchsian group, then we obtain the product structure

$$\overline{S}_j \cong \partial \overline{S}_j \times [0, 1],$$

with $CH(\Lambda(\Gamma_j)) = \partial \overline{S}_j \times \{1/2\}$. In the case that Γ_j *does* contain an elliptic element of order two whose axis is contained in $CH(\Lambda(\Gamma_j))$, then Σ_j is a neighborhood of a geodesic plane, and this order two element of $\Gamma_j = \text{Stab}_\Gamma(\Sigma_j)$ exchanges the two halves of this neighborhood. Such an order two element makes $CC(S_j)$ non-orientable as a 2-orbifold, and 1-sided as a 2-suborbifold. In particular, a regular

neighborhood of $CC(S_j)$ in this case will have boundary a 2-sided 2-suborbifold $F \subset Q$ which doubly covers $CC(S_j)$, and we have a homeomorphism $\partial\bar{S}_j \times \{t\} \cong F$ for any $t \in [0, 1)$. Finally, when $CC(S_j)$ is 3-dimensional and F is the hyperbolic 2-orbifold in $\partial CC(S_j)$ which separates the convex core of S_j from $f(\mathcal{T})$, we have the homeomorphism $\partial\bar{S}_j \times \{t\} \cong F$ for any $t \in [0, 1)$.

Consider the projection $r_\partial: \partial\bar{\Sigma}_j \rightarrow CH(\Lambda(\Gamma_j))$ given by restricting r . Because r is distance decreasing [13, Lemma 1.3.4], r_∂ is area decreasing. So the image $\Phi \circ r_\partial(\partial\bar{\Sigma}_j)$ in Q is a collection of hyperbolic 2-orbifolds with total area less than twice the area of \mathcal{T} . Define R_j to be the image under the projection Φ of the product region between $\partial\bar{\Sigma}_j$ and $r_\partial(\partial\bar{\Sigma}_j)$, the possibilities for which were described in the previous paragraph. So we have

$$R_j = S_j - CC(S_j).$$

(For reasons which will be obvious in Section 4, this product region will be called a *room*). In Section 5, we will see that it is impossible for $\Phi \circ r_\partial(\partial\bar{\Sigma}_j)$ to be anything other than a collection of hyperbolic turnovers or mirrored hyperbolic triangles (Corollary 5.5). Thus R_j is homeomorphically a collection of turnovers crossed with an interval. We will also see, in Section 4, that in this case we have the volume bound

$$\text{Vol}(R_j) < \frac{H}{2} \text{Area}(\partial\bar{S}_j),$$

where $H = 1.199678\dots$ is the positive solution of $x = \coth x$. We summarize these remarks as follows:

Lemma 3.6. *Suppose S_j is a component of $Q - f(\mathcal{T})$ and $\pi_1(S_j)$ is non-elementary. Then the region R_j between $\partial\bar{S}_j$ and $CC(S_j)$ is a product of a turnover with an open interval, and $\text{Vol}(R_j) < \frac{H}{2} \text{Area}(\partial\bar{S}_j)$, for H described above.*

3.5. Case: $|\Lambda(\Gamma_j)| = 2$. Let $\Lambda(\Gamma_j) = \{x, y\} \subset S_\infty^2$. Then the convex hull is the geodesic \tilde{l} connecting x and y , and the support planes for this convex hull consist of all planes containing \tilde{l} . It follows that Σ_j is a neighborhood of \tilde{l} , and Γ_j stabilizes this neighborhood, acting by translations or as an infinite dihedral group along \tilde{l} , and possibly also by rotations around \tilde{l} . Figure 5 provides a schematic/lower-dimensional illustration of a collection of Γ_j -invariant planes in \mathbb{H}^3 cutting out a neighborhood of a stabilized geodesic.




FIGURE 5. Schematic diagram for the two-point limit set case

Thus $f(\mathcal{T})$ cuts out a neighborhood of the orbifold geodesic $l = \Phi(\tilde{l})$ in Q . In this circumstance, we can put a complete hyperbolic structure on $Q_l = Q - l$ [18, Theorem 1.2.1], and so we are left with a new hyperbolic 3-orbifold, with a new cusp C , which again contains a new hyperbolic turnover $g(\mathcal{T})$. This new cusp lies in a component of $Q_l - g(\mathcal{T})$ whose fundamental group has limit set consisting either of 1 point or infinitely many points. This cusp admits deformations (for example, simply fill the cusp back in to obtain Q), so in the former case we would have that $g(\mathcal{T})$ cuts out a non-rigid cusp neighborhood in Q_l . We have remarked in Section 3.3 that such a phenomenon does not occur for an immersed turnover (see Lemma 3.5), and so we must be in the latter case. Furthermore, we have remarked that if the fundamental group of this component has infinite limit set, then Corollary 5.5 tells us its convex core must be a 3-dimensional orbifold with boundary a collection of embedded totally geodesic hyperbolic turnovers \mathcal{S} separating it from the rest of Q_l .

We observe that any such turnover boundary component will remain π_1 -injective under the hyperbolic Dehn filling on the cusp C which yields our original orbifold Q . This is by the Equivariant Loop Theorem [23], which implies that any embedded 2-orbifold in a good 3-orbifold which is not π_1 -injective has a compressing orbifold disk. If some elements of $\pi_1(\mathcal{S})$ were made trivial by filling C to retrieve Q , then there would have to be compressing orbifold disks to achieve this

for the image of \mathcal{S} in Q . But a turnover admits no compressing orbifold disks, because it admits no essential simple closed curves. We therefore have a non-empty collection of embedded totally geodesic turnovers in Q_l which separate the cusp C from the immersed turnover $g(\mathcal{T})$, and this collection survives any hyperbolic Dehn surgery performed on C . In particular, the collection survives if we fill C to retrieve Q . But now the collection \mathcal{S} geometrically isolates C from $g(\mathcal{T})$. This implies that the geometry of the immersion $g: \mathcal{T} \rightarrow Q_l$ does not change under fillings of C . In particular, there is a hyperbolic 3-orbifold V which has some totally geodesic boundary components, which contains a totally geodesic immersion $h(\mathcal{T})$ of the turnover \mathcal{T} , and for which the pair $(V, h(\mathcal{T}))$ isometrically embeds in both $(Q, f(\mathcal{T}))$ and $(Q_l, g(\mathcal{T}))$. But this is a contradiction, because we began by assuming our immersed turnover $f: \mathcal{T} \rightarrow Q$ cut out tubular neighborhood of an orbifold geodesic, which cannot contain any totally geodesic hyperbolic turnovers. This contradiction rules out the two-point limit set case:

Lemma 3.7. *If S_j is a component of $Q - f(\mathcal{T})$, then $\pi_1(S_j)$ can not be elementary with a two-point limit set.*

4. THE “GEODESIC FLOOR” ISOPERIMETRIC INEQUALITY

In order to prove Lemma 3.6, we need a tool for bounding the volume between S_j and $CC(S_j)$, when $\pi_1(S_j)$ is non-elementary and either $CC(S_j)$ is 2-dimensional and totally geodesic or $CC(S_j)$ is 3-dimensional and $\partial CC(S_j)$ is totally geodesic. This will be provided by the next theorem. First, we need some definitions and notation.

Let F be a measurable subset of finite area in a geodesic plane of \mathbb{H}^3 . Choose one closed half-space $H = \mathbb{H}_+^3 \subset \mathbb{H}^3$ of the plane containing F . For any $z \in F$, let l_z denote the geodesic ray perpendicular to F that begins at z and points into H . A *graph* over F is the image of a function $g: F \rightarrow H$ which maps $z \in F$ to the point on l_z at (hyperbolic) distance $\varphi(z)$ from z , where $\varphi: F \rightarrow \mathbb{R}_{\geq 0}$ is a non-negative function. Let C be a graph over F and $R := R(F, C)$ the region trapped between F and C . We will call F , C and R the *floor*, the *ceiling* and the *room*, respectively.

Let $S(r)$ denote the graph of constant height r over F , and let $B(r)$ denote the region trapped between $S(r)$ and F . We will refer to such an S (respectively, B) as a *nice ceiling* (respectively, *nice room*).

Theorem 4.1. (*Geodesic Floor Isoperimetric Inequality*) *Let C be the ceiling for the graph $g: F \rightarrow H$ of a function $\varphi: F \rightarrow \mathbb{R}_{\geq 0}$ which,*

for simplicity, is assumed to be differentiable almost everywhere. Let $R(F, C)$ be the associated room. Let B be the nice room over F with $\text{Vol}(B) = \text{Vol}(R)$, and let S be the ceiling of B . Then $\text{Area}(C) \geq \text{Area}(S)$.

Proof of 4.1. We take the metric on \mathbb{H}^3 given by $dh^2 + \cosh^2 h(dr^2 + \sinh^2 r d\theta^2)$, where r and θ describe polar coordinates in the plane containing F and h is the positive hyperbolic distance into H from F . We also denote the area form $\sinh r dr d\theta$ on the plane containing F by dA . Let $g: F \rightarrow H$ be given, and C the image of g . The volume of the room R is given by

$$(4.1) \quad V := \text{Vol}(R) = \iint_F \int_0^{g(r, \theta)} \cosh^2 h dh dA \\ = \iint_F \frac{1}{4} (\sinh 2g(r, \theta) + 2g(r, \theta)) dA,$$

and since the volume of the nice room B is the same, we have the constant H defined implicitly by

$$(4.2) \quad \text{Vol}(B) = \iint_F \frac{1}{4} (\sinh 2H + 2H) dA = V.$$

A calculation yields the following formula and lower bound for the area of C

$$(4.3) \quad \text{Area}(C) = \iint_F \cosh g(r, \theta) \sqrt{(g_r^2(r, \theta) + \cosh^2 g(r, \theta)) \sinh^2 r + g_\theta^2(r, \theta)} dr d\theta \\ \geq \iint_F \cosh g(r, \theta) \sqrt{\cosh^2 g(r, \theta) \sinh^2 r} dr d\theta \\ = \iint_F \cosh^2 g(r, \theta) dA.$$

The constant height H ceiling S of the nice room B has area

$$(4.4) \quad \text{Area}(S) = A_F \cosh^2 H = A_F (\sinh^2 H + 1),$$

where $A_F = \text{Area}(F)$. Because H is a constant, (4.2) can be rewritten as

$$(4.5) \quad \sinh 2H + 2H = 4V/A_F.$$

We therefore have

$$\frac{\frac{4V}{A_F} - 2H}{2 \cosh H} = \frac{\sinh 2H}{2 \cosh H} = \frac{2 \sinh H \cosh H}{2 \cosh H} = \sinh H,$$

which implies

$$A_F(\sinh^2 H + 1) = \frac{A_F^2 \cosh^2 H + (2V - H A_F)^2}{A_F \cosh^2 H}.$$

Substituting (4.4) into the above yields

$$\text{Area}(S) = \frac{A_F \text{Area}(S) + (2V - H A_F)^2}{\text{Area}(S)},$$

or, by the quadratic formula,

$$(4.6) \quad \text{Area}(S) = \frac{A_F + \sqrt{A_F^2 + 4(2V - H A_F)^2}}{2}.$$

We would like to show that the right-hand side of (4.6) is always less than or equal to $\iint_F \cosh^2 g \, dA$, and therefore, by (4.3), always less than or equal to $\text{Area}(C)$. We begin by observing that

$$(4.7) \quad \iint_F H \, dA \geq \iint_F g \, dA,$$

for otherwise equations (4.1) and (4.2) would imply

$$\frac{1}{A_F} \iint_F \sinh 2g \, dA < \sinh 2H < \sinh 2 \left(\frac{1}{A_F} \iint_F g \, dA \right),$$

which would violate Jensen's inequality, because $x \mapsto \sinh 2x$ is a convex function on $[0, \infty)$. Multiplying both sides of (4.7) by 2 and adding $\iint_F \sinh 2g \, dA$ to both sides, we can apply (4.1) to obtain

$$\iint_F 2 \sinh g \cosh g \, dA = \iint_F \sinh 2g \, dA \geq 4V - 2H A_F.$$

Now applying the Cauchy–Schwarz inequality to the left-most integral, we obtain

$$\sqrt{\iint_F 2 \sinh^2 g \, dA} \sqrt{\iint_F 2 \cosh^2 g \, dA} \geq 4V - 2H A_F.$$

The right-hand side of the above inequality is equal to $A_F \sinh 2H$, which is non-negative. We may therefore square both sides and apply a hyperbolic trigonometric identity to obtain

$$\left(\iint_F \cosh 2g - 1 \, dA \right) \left(\iint_F \cosh 2g + 1 \, dA \right) \geq 4(2V - HA_F)^2,$$

and expanding the left-hand side of the above gives

$$\left(\iint_F \cosh 2g \, dA \right)^2 - \left(\iint_F dA \right)^2 \geq 4(2V - HA_F)^2.$$

We rewrite this as

$$\left(\iint_F \cosh 2g + 1 - 1 \, dA \right)^2 \geq A_F^2 + 4(2V - HA_F)^2,$$

which is equivalent to

$$\left(2 \iint_F \cosh^2 g \, dA - A_F \right)^2 \geq A_F^2 + 4(2V - HA_F)^2.$$

Taking square roots, adding A_F to both sides, dividing by 2, and applying (4.3) and (4.6) to the result finally yields

$$\begin{aligned} (4.8) \quad \text{Area}(C) &\geq \iint_F \cosh^2 g \, dA \\ &\geq \frac{A_F + \sqrt{A_F^2 + 4(2V - HA_F)^2}}{2} \\ &= \text{Area}(S), \end{aligned}$$

which is what we wanted to show. 4.1

In Section 10, we will see that if B is a nice room (over any floor) and S is its ceiling, then

$$\frac{\text{Area}(S)}{\text{Vol}(B)} > \frac{2}{H},$$

where $H = 1.199678\dots$ is the positive solution of $x = \coth x$. As a result, we have the following corollary. It implies Lemma 3.6.

Corollary 4.2. *Let $R(F, C)$ be as in Theorem 4.1. Then $\text{Vol}(R) < \frac{H}{2} \text{Area}(C)$.*

5. A HYPERBOLIC “NEIGHBORHOOD” OF AN IMMERSSED TURNOVER

The goal of this section is to prove the first part of item (5) from Theorem 1.1, i.e., that an immersed turnover $f: \mathcal{T} \rightarrow Q$ in a complete orientable hyperbolic 3-orbifold is contained in a small hyperbolic 3-suborbifold with (possibly empty) totally geodesic boundary. Recall that a compact 3-orbifold is *small* if it is irreducible, contains no essential orientable 2-suborbifold and has (possibly empty) boundary consisting only of turnovers. The uniformization theorem for small 3-orbifolds [3] says that a small 3-orbifold has a geometric structure. We continue to denote the covering projection by $\Phi: \mathbb{H}^3 \rightarrow Q$, a connected component of $\mathbb{H}^3 - \Phi^{-1}(f(\mathcal{T}))$ by Σ_j , and the stabilizer of such a component by Γ_j .

Let $\mathcal{N}(W)$ denote a closed regular neighborhood of a subset $W \subset Q$. Consider the union of $\mathcal{N}(f(\mathcal{T}))$ and all of the components of $Q - f(\mathcal{T})$ whose fundamental group has finite limit set. From this union, remove small embedded open cusp neighborhoods. Call the result of this construction N . It is a compact 3-orbifold, and ∂N is non-empty if and only if there is a component of $Q - f(\mathcal{T})$ with infinite fundamental group that is not a neighborhood of an orbifold geodesic.

Theorem 5.1. *N is a small 3-orbifold. Moreover, the geometric structure on N is hyperbolic.*

Proof of 5.1. We begin by showing that N is irreducible. Consider the covering space \tilde{N} of N consisting of a connected component of $\Phi^{-1}(N)$. It is a noncompact manifold with boundary. We will show that \tilde{N} is irreducible and therefore, since \tilde{N} covers N , that N is irreducible as well [4, Remark following Theorem 3.23].

Let Y be an embedded sphere in \tilde{N} . We will show that Y bounds a ball. To do this, it is helpful to understand what \tilde{N} looks like. Of course, we have that one component X of $\Phi^{-1}(f(\mathcal{T}))$ will be contained in \tilde{N} . The complement $\mathbb{H}^3 - \Phi^{-1}(f(\mathcal{T}))$ consists of precompact open balls and non-precompact open balls (this is best visualized using the projective ball model for \mathbb{H}^3). The stabilizer for any one of these precompact open balls must be finite because the covering action on \mathbb{H}^3 is discrete, and therefore each such ball covers a component of

$Q - f(\mathcal{T})$ with finite fundamental group. In particular, each open ball in $\mathbb{H}^3 - X$ with finite volume is added to X as part of the construction of \tilde{N} .

We remark that we have not yet ruled out the possibility that $f(\mathcal{T})$ cuts out a solid torus or pillow from Q (Lemma 3.7), for this result depends on both Lemma 3.6 and the theorem we are in the middle of proving. As a consequence, we must include this possibility in the present analysis. We therefore also must add to X (as part of the construction of \tilde{N}) any non-precompact ball whose closure meets X and which covers a component of $Q - f(\mathcal{T})$ that looks like a solid torus or pillow. These non-precompact balls just look like infinite solid cylinders.

If a component of $\mathbb{H}^3 - X$ whose closure meets X is isotopic to a horoball centered at some $z \in S_\infty^2$, then we also add this component minus a smaller horoball around z to X as part of \tilde{N} . These pieces of \tilde{N} look like “horo-slabs,” and are the truncated developments in \mathbb{H}^3 of components of $Q - f(\mathcal{T})$ that look like cusp neighborhoods.

To finally complete the construction of \tilde{N} , we consider a component of $Q - f(\mathcal{T})$ whose fundamental group has infinite limit set Λ . Recall, from the discussion in the first paragraph of Section 3.4, that there is a product region between X and $\partial CH(\Lambda)$ (or between X and $CH(\Lambda)$, if the convex hull is 2-dimensional). We add on the “half” of this product region which meets X to complete the construction of \tilde{N} . It follows from this construction that $\partial\tilde{N}$ consists solely of embedded noncompact surfaces in \mathbb{H}^3 (corresponding either to horospheres or to the “halfway mark” for the flow onto $CH(\Lambda(\Gamma_j))$ for a component Σ_j of $\mathbb{H}^3 - \Phi^{-1}(f(\mathcal{T}))$ whose stabilizer has infinite limit set).

Returning to our embedded sphere Y , we observe that Y bounds a 3-ball B in \mathbb{H}^3 . If B is contained in \tilde{N} , then we are done. If $B \not\subset \tilde{N}$, then some interior point $z \in B$ is not contained in \tilde{N} . There are but two possibilities for the location of z : it is contained in a horoball region of $\mathbb{H}^3 - \tilde{N}$ or it is contained in one of the regions Σ_j of $\mathbb{H}^3 - \Phi^{-1}(f(\mathcal{T}))$ whose stabilizer has infinite limit set. A properly embedded noncompact surface separates \tilde{N} from any such region. Since $Y \subset \tilde{N}$, we therefore have that a properly embedded noncompact surface is contained in B . This contradiction implies that we must have $B \subset \tilde{N}$, and we conclude that N is irreducible.

We must show that N contains no essential 2-orbifolds. We first prove two lemmata.

Lemma 5.2. *N is atoroidal.*

Proof of 5.2. We have just seen that N is irreducible, so in order to show that N is atoroidal it is enough to show that an embedded torus, pillowcase, or Euclidean turnover in N which is incompressible is also ∂ -parallel. So assume that $S \subset N$ is such a 2-suborbifold, and suppose S is not parallel to a boundary component of N . Observe that as a 2-suborbifold of the orbifold $Q \supset N$, S must either compress or be parallel into a cusp of Q , because Q is atoroidal.

In the latter case, the product region determined by the isotopy of S into a cusp neighborhood of Q exhibits the isotopy of S into the boundary of N , because this cusp neighborhood must have been added to $Q - f(\mathcal{T})$ in the construction of N . The fact that this isotopy remains in N follows as in the irreducibility argument above, for the lift of the isotopy to the universal cover is a product region between two horospheres centered at the same point on S_∞^2 . If the isotopy of S does not remain in N , then there is a properly embedded noncompact surface (which is not a horosphere) in this product region. This is impossible.

If, on the other hand, S compresses in Q , then S must be a torus or pillowcase (S cannot be a Euclidean turnover because an embedded circle on a turnover always bounds an orbifold disk in the turnover). In this case S either bounds a solid orbifold torus or solid orbifold pillow, or S is contained in a Haken ball in Q (e.g. [4, Proposition 3.14]). If S is contained in a Haken ball in Q , then it is homotopically trivial, and again consideration of the construction of N allows us to conclude that S compresses in N . If S bounds a solid orbifold torus/pillow V , and if V is not contained in N , then V must contain a component of $Q - N$. Again, as in the irreducibility argument, we have a contradiction. This proves the lemma. 5.2

Lemma 5.3. *If $g: S \rightarrow N$ is an embedded hyperbolic turnover, then S is ∂ -parallel.*

Proof of 5.3. In the orbifold Q , S is isotopic to a totally geodesic embedding or to the double cover of mirrored triangle, which we may as well call S in either case. Consider any curve $\mathcal{C} \subset S \cap f(\mathcal{T})$. Because \mathcal{C} is in the intersection of two totally geodesic orbifolds, it is a geodesic. As a result, \mathcal{C} cannot be simple, because a geodesic on a hyperbolic turnover must intersect itself. However, we observe that $f^{-1}(\mathcal{C})$ must

be embedded in \mathcal{T} , because otherwise we would have an element of $\pi_1(\mathcal{T})$, corresponding to the curve $f^{-1}(\mathcal{C})$, taking one lift \tilde{S}_1 of S to another lift \tilde{S}_2 with $\tilde{S}_1 \neq \tilde{S}_2$ but $\tilde{S}_1 \cap \tilde{S}_2 \neq \emptyset$. Because S is embedded, this is not possible and we conclude that $f^{-1}(\mathcal{C})$ is embedded in \mathcal{T} . But an embedded curve on a turnover bounds an orbifold disk, and this implies that the geodesic \mathcal{C} is contractible, which is a contradiction. This contradiction also holds in N because S was made totally geodesic in Q by an isotopy. In particular, S is disjoint from $f(\mathcal{T})$ in N , and since we know what the complementary pieces of $N - f(\mathcal{T})$ look like, the conclusion follows. 5.3

We can use these lemmata to show that N has a hyperbolic structure. To do this, we use the turnover splitting of a 3-orbifold [4, Theorem 4.8]. This says that a compact, irreducible, atoroidal 3-orbifold \mathcal{O} contains a maximal (possibly empty) collection \mathfrak{h} of essential, pairwise disjoint, non-parallel, hyperbolic turnovers which is unique up to isotopy and such that every component of \mathcal{O} split along \mathfrak{h} is either small or Haken. Lemma 5.2 shows that N is atoroidal, and Lemma 5.3 implies that the turnover splitting collection of N is empty. So N is either small or Haken. As we claim, it will turn out that N is small.

Now suppose that $g: S \rightarrow N$ is a proper embedding of a ∂ -incompressible, incompressible, orientable 2-orbifold in N . Make $g(S)$ transverse to $f(\mathcal{T})$ in N , and consider any curve $\mathcal{C} \subset f(\mathcal{T}) \cap g(S)$. As in the hyperbolic turnover case, we must have that $f^{-1}(\mathcal{C})$ is simple in \mathcal{T} , for otherwise a multiple point would give rise to a contradiction that g is an embedding. In particular, $f^{-1}(\mathcal{C})$ bounds an orbifold disk $D_{\mathcal{T}}$ in \mathcal{T} , and therefore \mathcal{C} is contractible along $f(D_{\mathcal{T}})$ in N . Now since $g(S)$ is essential, the map g is π_1 -injective, and so \mathcal{C} bounds a (possibly non-embedded) disk D_S in $g(S)$. By choosing \mathcal{C} which has innermost preimage in \mathcal{T} , we can use the irreducibility of N to obtain a map of an orbifold ball B into N with boundary $f(D_{\mathcal{T}}) \cup D_S$. We may then use B to remove \mathcal{C} from $f(\mathcal{T}) \cap g(S)$ and thus decrease the number of curves of intersection of the two orbifolds. After a finite number of steps, we see that $g(S)$ can be made disjoint from $f(\mathcal{T})$. But we know what the components of $N - f(\mathcal{T})$ look like. Namely, these components are either ∂ -parallel product regions, orbifold balls, solid tori, or solid pillows. We conclude that $g(S)$ is ∂ -parallel. This proves that N contains no essential, orientable 2-suborbifolds.

As we observed above, N is either small or Haken. Because N is irreducible and contains no embedded, orientable, essential 2-suborbifold, consideration of the construction of N shows that it is small *exactly* when there is no component $\mathcal{Q} \subset Q - f(\mathcal{T})$ satisfying any of the following conditions:

- (1) \mathcal{Q} is a non-rigid cusp neighborhood, or
- (2) $CC(\mathcal{Q})$ is a non-rigid hyperbolic 2-orbifold, or
- (3) $\partial CC(\mathcal{Q})$ contains a non-rigid hyperbolic 2-orbifold which separates $CC(\mathcal{Q})$ from $f(\mathcal{T})$.

Any such \mathcal{Q} will produce a non-turnover boundary component in the construction of N , and thus prevent N from being small. So to show that N is small, we suppose for the sake of contradiction that there is such a \mathcal{Q} . In this case, we see that N satisfies the following theorem of Dunbar [11] (note that N contains no bad 2-suborbifold, i.e., N is “abad,” because $Q \supset N$ can contain no such 2-suborbifold):

Theorem 5.4. *Let \mathcal{O} be a smooth, compact, connected, irreducible, abad, orientable 3-orbifold in which every non-spherical turnover is boundary-parallel. If $\partial\mathcal{O}$ has a component which is not a turnover, then \mathcal{O} has a “strong hierarchy,” i.e., \mathcal{O} can be decomposed into orbifold balls and thick turnovers by repeated cutting along 2-sided, essential 2-suborbifolds.*

But we have already seen that N contains no 2-suborbifolds for such a hierarchy. Since N is neither an orbifold ball nor a thick turnover (the latter because \mathcal{T} is not embedded), we conclude that there can be no such component \mathcal{Q} as above. Therefore, N is small, and either $\text{int}(N) \cong Q$ or N is obtained from a component Q' of Q cut along a collection of totally geodesic turnovers and 1-sided triangles with mirrored sides and possibly also by truncating rigid cusps. (This component Q' is what we refer to in (5) of Theorem 1.1.) In particular, N is hyperbolic. This proves the theorem. 5.1

Corollary 5.5. *Every cusp neighborhood of Q which is cut out by $f(\mathcal{T})$ must be rigid. If \mathcal{Q} is a component of $Q - f(\mathcal{T})$ with non-elementary fundamental group, then $CC(\mathcal{Q})$ is either a totally geodesic hyperbolic turnover, a totally geodesic 1-sided triangle with mirrored sides, or a 3-dimensional suborbifold with boundary a collection of hyperbolic turnovers which separate $CC(\mathcal{Q})$ from $f(\mathcal{T})$.*

6. BOUNDING THE TYPE OF AN ELLIPTIC IMMERSION

In this section, we will introduce the tools necessary to prove items (3) and (4) of Theorem 1.1, which refer to the orders of the cone points or dihedral angles of the collection of splitting totally geodesic orbifolds which form the boundary of the “turnover core” Q' . The collection of tools we need comes from the work of Gehring and Martin [14], [21].

If l_1 and l_2 are geodesics in \mathbb{H}^3 , then denote by $\rho(l_1, l_2)$ the length of the unique common perpendicular segment between l_1 and l_2 , or zero if l_1 and l_2 intersect in $\mathbb{H}^3 \cup S_\infty^2$. Let n and m be positive integers with $n \geq \max\{3, m\}$. Define $c(n, m)$ by

$$(6.1) \quad c(n, m) = \begin{cases} \sqrt{2 \cos(2\pi/n) - 1}/2 & \text{if } n \geq 7, \\ \cos(\pi/m)/2 & \text{if } n = 6 \text{ and } m \geq 3, \\ 1/\sqrt{8} & \text{if } n = 6 \text{ and } m = 2, \\ \sqrt{(\sqrt{5} - 1)/16} & \text{if } n = 5, \\ \sqrt{(\sqrt{3} - 1)/8} & \text{if } n = 4, \\ \sqrt{(\sqrt{5} - 2)/8} & \text{if } n = 3. \end{cases}$$

Now let g and h be elliptic isometries of order n and m , respectively, which generate a Kleinian group. Gehring and Martin show that either $\rho(\text{axis}(g), \text{axis}(h)) = 0$ or that the following inequality holds [14, Theorem 6.19]:

$$(6.2) \quad \rho(\text{axis}(g), \text{axis}(h)) \geq \sinh^{-1} \left(\frac{c(n, m)}{\sin(\pi/n) \sin(\pi/m)} \right).$$

Denote the right-hand side of (6.2) by $\delta(n, m)$. Observe that it provides a lower bound for the “axial distance” between any two elliptic types, i.e., whenever the axes of any two elliptic isometries of orders n and m come closer than $\delta(n, m)$, it is necessary that the group generated by these isometries be elementary. When $n \geq m \geq 7$, any two elliptic elements of orders n and m which generate a discrete group cannot have axes which meet in $\mathbb{H}^3 \cup S_\infty^2$, so $\delta(n, m) > 0$ always gives a lower bound for axial distance in this case. We will be most interested in the case $n = m \geq 7$, for which we have the strictly increasing function [14, Example 8.11], [21, Theorem 2.2]:

$$(6.3) \quad \delta(n, n) = 2 \cosh^{-1} \left(\frac{1}{2 \sin(\pi/n)} \right).$$

We will show in Section 10 that the injectivity radius at any point of a turnover is bounded above by the maximal radius r_{max} of an embedded disk in a hyperbolic thrice-punctured sphere (Proposition 10.1). We also calculate the value of r_{max} :

$$(6.4) \quad r_{max} = \ln \left(\frac{2 + \sqrt{7}}{\sqrt{3}} \right) = 0.986647\dots$$

Recall that Π_T denotes the unique plane stabilized by the turnover subgroup $T = T(p, q, r) \cong \pi_1(\mathcal{T}) \leq \Gamma \cong \pi_1(Q)$, and $\Phi: \mathbb{H}^3 \rightarrow Q$ the covering projection. Suppose that $\gamma \in \Gamma$ is an elliptic element of order $n \geq 7$, with $\text{axis}(\gamma) \cap \Pi_T \neq \emptyset$. Then $\text{axis}(\gamma)$ and Π_T are either perpendicular or not. Suppose they are not perpendicular. We cannot have $\text{axis}(\gamma) \subset \Pi_T$, because then some T -translate of $\text{axis}(\gamma)$ would intersect $\text{axis}(\gamma)$, and this would violate discreteness because γ has order at least 7. So the intersection is a point (corresponding to a smooth point of \mathcal{T} , but a singular point of Q), and the action of γ provides the local model for part of the immersion $f: \mathcal{T} \rightarrow Q$ (i.e., $f(\mathcal{T}) \subset Q$ near $\Phi(\text{axis}(\gamma) \cap \Pi_T)$ looks like a plane “spun” around an oblique line). Since $n \geq 7$, we must have that $\rho(\text{axis}(\gamma), \text{axis}(\alpha\gamma)) \geq \delta(n, n)$ for all $\alpha \in \Gamma - \text{Stab}(\text{axis}(\gamma))$, that is, $\text{axis}(\gamma)$ has a tubular neighborhood of radius $\delta(n, n)/2$ which is disjoint from any translate by a deck transformation which does not stabilize it. This is equivalent to $\Phi(\text{axis}(\gamma))$ having an embedded (orbifold) tubular neighborhood of radius $\delta(n, n)/2$ in Q . But once we have $\delta(n, n)/2 > r_{max}$, there will be an element $\alpha \in T$ for which the tubular neighborhoods of radius $\delta(n, n)/2$ around $\text{axis}(\gamma)$ and $\text{axis}(\alpha\gamma)$ will not be disjoint. We therefore conclude that $\delta(n, n)/2$ must be less than r_{max} , and an easy calculation shows that this occurs only for $n \leq 9$.

Now suppose that $\text{axis}(\gamma)$ meets Π_T in a right angle. Then the group $\langle \gamma, T \rangle$ stabilizes Π_T , and so is an orientation-preserving Fuchsian group. It is known that turnover subgroups are maximal among orientation-preserving Fuchsian groups, so in fact $\langle \gamma, T \rangle$ is a turnover subgroup of $\pi_1(Q)$. There are two ways in which this can occur. One is that $\langle \gamma, T \rangle = T$, in which case we have $\gamma \in T$, and so $\text{axis}(\gamma)$ is contained in the axis of a (maximal) elliptic element of order p, q , or r .

$T(l, m, n)$	$\geq T(p, q, r)$	Index	Normal
$(3, 3, t)$	(t, t, t)	3	Yes
$(2, 3, 2t)$	(t, t, t)	6	Yes
$(2, s, 2t)$	(s, s, t)	2	Yes
$(2, 3, 7)$	$(7, 7, 7)$	24	No
$(2, 3, 7)$	$(2, 7, 7)$	9	No
$(2, 3, 7)$	$(3, 3, 7)$	8	No
$(2, 3, 8)$	$(4, 8, 8)$	12	No
$(2, 3, 8)$	$(3, 8, 8)$	10	No
$(2, 3, 9)$	$(9, 9, 9)$	12	No
$(2, 4, 5)$	$(4, 4, 5)$	6	No
$(2, 3, 4t)$	$(t, 4t, 4t)$	6	No
$(2, 4, 2t)$	$(t, 2t, 2t)$	4	No
$(2, 3, 3t)$	$(3, t, 3t)$	4	No
$(2, 3, 2t)$	$(2, t, 2t)$	3	No

TABLE 1. Turnover Supergroups and Subgroups

The other way that $\langle \gamma, T \rangle$ can be a turnover subgroup is that T is contained in some turnover *supersubgroup* $T'(l, m, n) \leq \pi_1(Q)$, or equivalently, when $f(\mathcal{T})$ covers a smaller immersed turnover $f': \mathcal{T}' \rightarrow Q$ (recall that the hypothesis of the main theorem is that $f(\mathcal{T})$ does not cover an *embedded* turnover). If this happens, then we must know what types of turnovers can be covered by \mathcal{T} . The subgroups and supergroups of a turnover group can be determined from Table 1, whose data is collected from Singerman [29]. The data in the first column of the table gives the turnover groups that contain turnover subgroups. The second column gives the turnover subgroups so contained. Any turnover group not listed in the second column is maximal. By analyzing the table, it is readily seen that the order of γ must live in the set

$$(6.5) \quad \{2, 3, 4, 5, 7, 8, 9, p, q, r, 2p, 2q, 2r\}.$$

Now if \mathcal{T}_i is a turnover or triangle in the splitting collection for Q' as in the main theorem, then the cone points or dihedral angles of \mathcal{T}_i arise from axes of elliptic elements such as γ above, that is, elliptic elements whose axes intersect Π_T in a single point. This is because such an elliptic axis projects directly to a cone point (or dihedral

point) in $\partial Q'$. If the cone points of the immersed turnover have orders p, q , and r , then by our discussion above, it is possible for the orders of the singular points of any embedded turnovers or triangles in the complement of $f(\mathcal{T})$ to take values in the set

$$(6.6) \quad \{2, 3, \dots, 9, p, q, r, 2p, 2q, 2r\}.$$

Items (3) and (4) of the main theorem follow.

7. PROOF OF THE MAIN THEOREM

In this section, we will put everything together and prove the main theorem. Let $f: \mathcal{T} \rightarrow Q$ be a totally geodesic immersion of a hyperbolic turnover $\mathcal{T} = \mathcal{T}(p, q, r)$ in a complete, orientable, hyperbolic 3-orbifold Q . Recall that $p, q, r \in \mathbb{Z}$ satisfy $2 \leq p \leq q \leq r$ and $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} < 1$. Assume, as in Theorem 1.1, that $f: \mathcal{T} \rightarrow f(\mathcal{T})$ is not a covering of an embedded turnover or triangle in Q . Let $\{\mathcal{T}_i\}$ be the collection of hyperbolic turnovers and triangles obtained from the convex cores of components of $Q - f(\mathcal{T})$ with non-elementary fundamental group. This collection is embedded, pairwise disjoint, and each element of the collection is disjoint from $f(\mathcal{T})$, which is item (1) of the theorem. Let S_j denote a component of $Q - f(\mathcal{T})$. Recall that projection from $\partial \overline{S}_j$ onto its image in $CC(S_j)$ strictly decreases area. We have the bound

$$(7.1) \quad \bigcup_j \text{Area}(\partial \overline{S}_j) \leq 2\text{Area}(\mathcal{T}) = 4\pi \left(1 - \left(\frac{1}{p} + \frac{1}{q} + \frac{1}{r}\right)\right)$$

on the total area of the boundary of these complementary components. The reason we do not have equality above is that it may be that $f(\mathcal{T})$ covers a smaller immersed orbifold in Q (equivalently, $\pi_1(\mathcal{T})$ is contained in some larger Fuchsian subgroup). It follows that the upper bound that we can give on the surface area of the complement of the turnover can be improved the more we know about the Fuchsian groups G such that $\pi_1(\mathcal{T}) \leq G \leq \pi_1(Q)$. These remarks will be useful in Section 9.

When $\pi_1(S_j)$ is non-elementary, we also have a lower bound of $\frac{\pi}{21}$ for the area of the image of the projection of $\partial \overline{S}_j$ onto $\partial CC(S_j)$, coming from twice the area of the $(2, 3, 7)$ mirrored triangle, which is the minimal area hyperbolic 2-orbifold. These two observations imply not only that there is a *global* upper bound on the number of elements in $\{\mathcal{T}_i\}$, but also an upper bound (which is at least as strong as the

global bound) on this number depending only on p, q , and r . This is item (2) of the theorem.

Now let Q' be the component of $Q - \cup_i \mathcal{T}_i$ which contains $f(\mathcal{T})$. The path metric closure of Q' is homeomorphic to the 3-orbifold N constructed in Section 5. It follows that the metric closure of Q' is a small hyperbolic orbifold with a totally geodesic boundary component for every element of $\{\mathcal{T}_i\}$. In Section 3, we showed that we have upper bounds on the volume of each piece of $Q' - f(\mathcal{T})$ in terms of the area of $f(\mathcal{T})$. We will use the volume estimates (and other claims) provided by Lemmata 3.2, 3.5, 3.6, and 3.7 in the calculation below, depending on the dimension of the convex core of a component $S_j \subset Q - f(\mathcal{T})$, which can be 3, 2, 0 or -1 (the latter occurs in the parabolic case; Lemma 3.7 tells us that $\dim CC(S_j) \neq 1$). Recall that $H = 1.199678\dots$ is the positive solution of the equation $x = \coth x$. We have

(7.2)

$$\begin{aligned}
\text{Vol}(Q') &= \sum_j \text{Vol}(S_j - CC(S_j)) \\
&= \sum_{\dim CC(S_j) < 1} \text{Vol}(S_j) + \sum_{\dim CC(S_j) > 1} \text{Vol}(S_j - CC(S_j)) \\
&< \sum_{\dim CC(S_j) < 1} \frac{1}{2} \cdot \text{Area}(\partial \bar{S}_j) + \sum_{\dim CC(S_j) > 1} \frac{H}{2} \cdot \text{Area}(\partial \bar{S}_j) \\
&< \frac{H}{2} \cdot \sum_j \text{Area}(\partial \bar{S}_j) \leq \frac{H}{2} \cdot 2 \cdot \text{Area}(\mathcal{T}) = H \cdot \text{Area}(\mathcal{T}).
\end{aligned}$$

This is almost all of item (5) of the theorem. The last statement follows from the fact that we have no terms above with H when there are no boundary components for the orbifold Q' . This completes the proof of (5) in Theorem 1.1.

We still need to prove the finiteness result, that is, that for a given p, q , and r , there are only finitely many possibilities for Q' . We use Jørgensen's Theorem [30, Theorem 5.12.1], [12, Theorem 5.5]: Given an upper bound K for volume, there is a finite collection $\mathcal{O} = \{O_1, O_2, \dots, O_k\}$ of hyperbolic 3-orbifolds such that any hyperbolic 3-orbifold with volume $< K$ is obtained from some element of \mathcal{O} by hyperbolic Dehn surgery. We can apply this theorem to the double DQ' of Q' along its boundary.

Since infinitely many hyperbolic 3-orbifolds are obtained by hyperbolic Dehn surgery on the elements of \mathcal{O} , we need to prove that only finitely many of these can contain an immersed (p, q, r) turnover. To do this, it is sufficient to show that no surgery with “large” coefficients on any element of \mathcal{O} can yield an orbifold containing an immersed (p, q, r) turnover.

An embedded neighborhood of a non-rigid cusp in a hyperbolic 3-orbifold is homeomorphic to either a solid torus minus its core curve or an open solid pillow minus its core singular curve. Hyperbolic Dehn surgery on a cusped hyperbolic 3-orbifold is performed by removing an embedded open neighborhood of a cusp, choosing an isotopy class of a closed curve c in the resulting boundary component V , and gluing either a solid torus or solid pillow (depending on the cusp) to V so that c is attached to the unique isotopy class of a meridian curve in the solid torus or pillow. The surgery has “large” coefficients if the isotopy class of c has slope $\frac{s}{t}$ (in the universal cover of V) and if $|s| + |t|$ is large. If hyperbolic Dehn surgery with large coefficients is performed on a cusped hyperbolic 3-orbifold, then the resulting core curve of the filled solid torus or pillow is either a short orbifold geodesic or an orbifold geodesic with a very small cone angle. Thurston’s Hyperbolic Dehn Surgery Theorem (e.g. [12, Theorem 5.3]) says that if one excludes a finite number of slopes from each cusp in a cusped hyperbolic 3-orbifold, then the hyperbolic Dehn surgeries on the remaining slopes all yield hyperbolic 3-orbifolds.

Suppose that DQ' is obtained from the collection \mathcal{O} by hyperbolic Dehn surgery, and that DQ' contains (two copies of) the immersed (p, q, r) turnover $f(\mathcal{T})$. Let $l \subset DQ'$ be the orbifold geodesic core of a solid torus or pillow coming from a filled cusp. Note that l corresponds to the axis A of a hyperbolic isometry of \mathbb{H}^3 which acts on a tubular neighborhood of A by either a translation and rotation or by a dihedral group and rotation. A result due to Meyerhoff [24] implies that l must have an embedded tubular neighborhood in DQ' and, furthermore, that the radius of such a neighborhood goes to infinity if either the length of l goes to zero or the cone angle on l goes to zero. Hence, a large hyperbolic Dehn surgery on an element of \mathcal{O} must yield a 3-orbifold with an embedded tube of large radius.

We have analyzed the components of $Q' - f(\mathcal{T})$ in detail, and seen that no such component supports a simple geodesic. So we must have $l \cap f(\mathcal{T}) \neq \emptyset$ or $l \cap S \neq \emptyset$, where S is an embedded turnover in DQ'

corresponding to a boundary component of Q' . Suppose first that l meets such an embedded turnover S , and suppose that l is contained in the singular locus of DQ' . Then the cone angle at l is 2π divided by an integer from the set $\{2, 3, \dots, 9, p, q, r, 2p, 2q, 2r\}$, by item (3) of the main theorem. As a consequence, the cone angle of such a filled cusp cannot be arbitrarily small. The other possibilities are either a non-singular geodesic which meets such an embedded turnover S , or else an orbifold geodesic for which $l \cap f(\mathcal{T}) \neq \emptyset$. But in both of these cases, we may apply the Meyerhoff bound mentioned above to conclude that l can be neither too short nor have a small cone angle, because a hyperbolic turnover has an upper bound on its injectivity radius, and a tube around l with radius larger than this maximal injectivity radius will not be embedded in DQ' . So l could not have come from a surgery on any element of \mathcal{O} with large coefficients. This proves the finiteness claim, and the main theorem.

8. FINE-TUNING THE SEARCH FOR IMMERSED TURNOVERS

In this section, we will generalize results of Miyamoto [25], with the goal (in Section 9) of limiting the types of embedded turnovers and triangles that can occur in the complement of a few specific immersed turnovers $\mathcal{T}(p, q, r)$. Let N be the hyperbolic 3-orbifold obtained by path metric completion of the splitting along the collection $\{\mathcal{T}_i\}$ of embedded turnovers and triangles from the main theorem. Then ∂N is a collection of hyperbolic turnovers, one for each \mathcal{T}_i .

We need some terminology. In the projective model of hyperbolic n -space \mathbb{H}^n , consider a linearly independent set of $n + 1$ points which lie either on the sphere at infinity S_∞^{n-1} or outside of the projective ball. If the line segment between each pair of these points intersects the interior of the projective ball, then the $n + 1$ points determine a *truncated n -simplex*. This is obtained by taking the infinite volume polyhedron in \mathbb{H}^n spanned by the points, and cutting off the infinite volume ends by the hyperplanes which are dual to the super-ideal vertices. A truncated simplex is *regular* if every edge between two of these truncating planes has the same length. If $r \geq 0$, then define $\rho_n(r)$ to be the ratio of the volume of a regular truncated n -simplex with edge length $2r$ to the $(n - 1)$ -volume of its truncated faces. The case of interest to us is dimension three, and we will write T_θ for a regular truncated 3-simplex whose non-truncated faces meet in the

angle θ , noting that θ and the edge length $2r$ of T_θ are related by

$$(8.1) \quad \cosh 2r = \frac{\cos \theta}{2 \cos \theta - 1}.$$

The truncated faces of regular truncated 3-simplex of angle θ are equilateral hyperbolic triangles with angle θ , and so if T_θ has edge length $2r$, then a result due to Miyamoto [25, Proposition 1.1] implies

$$(8.2) \quad \begin{aligned} \rho_3(r) &= \frac{\text{Vol}(T_\theta)}{4(\pi - 3\theta)} \\ &= \frac{1}{4(\pi - 3\theta)} \left(-8 \int_0^{\pi/4} \ln(2 \sin u) du - 3 \int_0^\theta \cosh^{-1} \left(\frac{\cos t}{2 \cos t - 1} \right) dt \right). \end{aligned}$$

A *return path* in a hyperbolic n -orbifold Q with boundary is a geodesic segment in Q which meets ∂Q perpendicularly at both of its end points.

All of Miyamoto's results are given for hyperbolic manifolds, but some of his arguments do not require any sort of assumptions of no torsion. In particular, the following result is true (cf. [25, Lemma 4.1]):

Proposition 8.1. *If a complete hyperbolic n -orbifold Q of finite volume with totally geodesic boundary has a lower bound $l \geq 0$ for the length of its return paths, then*

$$\text{Vol}(Q) \geq \rho_n \left(\frac{l}{2} \right) \text{Vol}(\partial Q).$$

Call a return path in Q *closed* if it begins and ends at the same point of ∂Q . Such a path must lift to the universal cover of Q as a geodesic segment which meets perpendicularly the axis of an order 2 elliptic element in $\pi_1(Q)$. With a little bit of careful analysis, we will prove the following result (cf. [25, Lemma 5.3]):

Proposition 8.2. *Let Q be a complete hyperbolic 3-orbifold with closed totally geodesic boundary. Then Q has a shortest return path γ , and there is a positive integer k such that the length of γ is at least the edge length of T_θ , where*

$$\theta = \frac{\pi}{3(1 - k\chi(\partial Q))}$$

if γ is closed and

$$\theta = \frac{\pi}{3(1 - \frac{k}{2}\chi(\partial Q))}$$

if γ is not closed. In both cases, $k > 1$ if and only if γ is contained in a singular axis of (maximal) order k in Q .

We can apply these results to our small hyperbolic 3-orbifold N . For a shortest return path γ in N and the corresponding truncated regular 3-simplex T_θ of side length $2r$ of Proposition 8.2, we have, by Theorem 1.1 and Proposition 8.1, the inequalities

$$(8.3) \quad \frac{1}{4(\pi - 3\theta)} \text{Vol}(T_\theta) \text{Area}(\partial N) = \rho_3 \left(\frac{2r}{2} \right) \text{Area}(\partial N) \\ \leq \text{Vol}(N) < H \cdot \text{Area}(\mathcal{T}(p, q, r)).$$

(Recall $H = 1.199678\dots$) It will turn out that we can use these inequalities to limit the number and kinds of turnovers which may appear in ∂N , by showing that, for certain types of turnovers in ∂N , the lower bound for $\text{Vol}(N)$ is greater than the upper bound. This will be done in the next section. To get there, we must first prove the proposition.

Proof of 8.2. The proof is identical to Miyamoto's proof, provided we take care to consider torsion. However, its elegance is worthy of replication, and we will find use for some of the ideas in Section 9.

Q has a shortest return path because its boundary is closed. Denote such a path by γ , and its length by l . Let γ_1 and γ_2 be lifts of γ to the universal cover $\tilde{Q} \subset \mathbb{H}^3$ (which is a convex region bounded by geodesic planes), such that γ_1 and γ_2 meet a common component P of $\partial\tilde{Q}$. The argument now diverges according to whether or not γ_1 and γ_2 are equal. There are two ways in which equality can occur. One is that γ is contained in the singular locus of Q , and the other is that γ passes perpendicularly through a singular axis of order 2 in Q . We will first suppose that $\gamma_1 \neq \gamma_2$, for all choices of lifts of γ . In particular, γ is not closed. Let P_1 and P_2 denote the other planes of $\partial\tilde{Q}$ which γ_1 and γ_2 meet, respectively, and let γ' denote the unique common perpendicular of P_1 and P_2 . Then the geodesics γ_1, γ_2 , and γ' are coplanar and form the alternating sides of an all-right hyperbolic hexagon. Let d be the distance of the side connecting the end points of γ_1 and γ_2 in P , and let l' denote the length of γ' . Observe that l'

is no less than l . So the all-right hexagon law of hyperbolic cosines implies

$$(8.4) \quad \cosh d = \frac{\cosh^2 l + \cosh l'}{\sinh^2 l} \geq \frac{\cosh l}{\cosh l - 1}.$$

Because the above inequality is valid for all choices of lifts of γ , there are two disjoint disks of radius $r = \frac{1}{2} \cosh^{-1}(\cosh l / (\cosh l - 1))$ contained in ∂Q and centered at the end points of γ .

We now may apply Böröczky's estimate for circle packings in spaces of constant curvature [5]. This says that, given a circle packing of radius r in a space of constant curvature, the density of each disk in its Dirichlet-Voronoi cell (that is, the collection of points lying nearer to the center of the disk than to any other disk in the packing) is at most the density of three mutually tangent disks of radius r in an equilateral triangle spanned by their centers. Let θ denote the angle of a hyperbolic equilateral triangle of side length $2r$. The area of ∂Q is $-2\pi\chi(\partial Q)$, and the area of a disk of radius r is $2\pi(\cosh r - 1)$, so we have the inequality

$$\frac{4\pi(\cosh r - 1)}{-2\pi\chi(\partial Q)} \leq \frac{3\theta(\cosh r - 1)}{\pi - 3\theta},$$

which implies that $\theta \geq \pi/(3(1 - \chi(\partial Q)/2))$. By the law of hyperbolic cosines, we obtain

$$\cosh 2r = \frac{\cos^2 \theta + \cos \theta}{\sin^2 \theta} \leq \frac{\cos(\pi/(3(1 - \chi(\partial Q)/2)))}{1 - \cos(\pi/(3(1 - \chi(\partial Q)/2)))}.$$

Now we have

$$\cosh l = \frac{\cosh 2r}{\cosh 2r - 1} \geq \frac{\cos(\pi/(3(1 - \chi(\partial Q)/2)))}{2\cos(\pi/(3(1 - \chi(\partial Q)/2))) - 1}.$$

But equation (8.1) implies that l is at least as large as the side length of a regular truncated 3-simplex with angle $\pi/(3(1 - \chi(\partial Q)/2))$. This proves the $k = 1$ case when γ is not closed.

The case when γ is closed but not contained in the singular locus follows in precisely the same manner, except that we now take torsion into account. So suppose that $\gamma_1 = \gamma_2$ and that γ is not contained in the singular locus of Q . Then γ must meet an order two axis of Q perpendicularly. If we consider another lift of γ which is distinct from γ_1 , then the same argument applies as above, except that we have only *one* disk packed into ∂Q rather than two. The same calculation yields

a bound for the length l of γ in terms of θ which is equal to the claim in the closed case with $k = 1$.

We are left with the case that γ is contained in the singular set of Q . Then there is an elliptic isometry in $\pi_1(Q)$ of some maximal order k whose axis contains γ_1 . Suppose first that $\gamma_1 = \gamma_2$ and that γ is closed. As in the last case, we pick a lift of γ which is distinct from γ_1 , and we conclude that we have one orbifold disk with cone point of order k packed in ∂Q . In order to get an estimate on the radius of this orbifold disk, we again appeal to Böröczky's estimate. Since our packing in ∂Q is with one orbifold disk, the Dirichlet-Voronoi cell of a disk in the packing obtained by lifting to the universal cover has area equal to $k\text{Area}(\partial Q)$ (because the disk is centered around the lift of an order k cone point). We therefore obtain, as above, the estimate for r as

$$\frac{2\pi(\cosh r - 1)}{k(-2\pi\chi(\partial Q))} \leq \frac{3\theta(\cosh r - 1)}{\pi - 3\theta}.$$

We now continue as in the proof of the first case, and the value for the bound on l in terms of θ so obtained is equal to the claim in the closed case with $k \neq 1$.

In the final case, we have two orbifold disks with cone points of order k packed in ∂Q (which is not necessarily connected). In the universal cover of ∂Q , we again have a packing by disks of some radius r . Looking at two lifts of the two orbifold disks packed into ∂Q , we have two associated Dirichlet-Voronoi cells V_1 and V_2 , each of whose area is k times the area of its image in ∂Q . We again apply Böröczky's result to this packing of two disks in two cells and obtain the estimate for r as

$$\frac{4\pi(\cosh r - 1)}{k(-2\pi\chi(\partial Q))} = \frac{4\pi(\cosh r - 1)}{\text{Area}(V_1 \cup V_2)} \leq \frac{3\theta(\cosh r - 1)}{\pi - 3\theta},$$

and the same analysis completes the proof. 8.2

9. APPLICATIONS AND EXAMPLES

In this section, we will apply the preceding results to some specific hyperbolic turnovers, as well as provide some examples which show that all of the hypotheses of the main theorem are necessary. In the last subsection, we will use the main theorem to correct an error in a result of Maclachlan [20], and to complete the classification of turnover

subgroups in the arithmetic cocompact hyperbolic tetrahedral reflection groups.

9.1. Immersed $(2, 4, 5)$ Turnovers. Let us analyze the case of an immersed $(2, 4, 5)$ hyperbolic turnover \mathcal{T} in a hyperbolic orbifold Q . Let N be the small orbifold obtained by splitting along embedded turnovers in Q , as in the main theorem. Then $\text{Vol}(N) < H \cdot \text{Area}(\mathcal{T}) = 0.376890\dots$, where $H = 1.199678\dots$ is given by Theorem 1.1. Using inequality (6.2), Table 1, and the fact that the hypotenuse of the $(2, 4, 5)$ hyperbolic triangle represents the greatest distance (approximately 0.842482, which is *less than* $\delta(n, 5)$ for all $n \geq 6$) separating two points on the triangle, it is easily seen that the only possibilities for the orders of cone points in ∂N are 2, 3, 4, 5. Because the projection from the immersed turnover onto any boundary pieces strictly decreases area, and because the total two-sided surface area of \mathcal{T} in Q is $\pi/5$, an easy calculation implies that the only possibilities for ∂N are a $(2, 4, 5)$ turnover or a $(3, 3, 4)$ turnover.

By considering cases, we can rule out a $(3, 3, 4)$ turnover boundary by using Proposition 8.2 and the inequalities (8.3). For example, because the $(3, 3, 4)$ turnover has only one order 4 cone point, we can conclude that, if the shortest return path in N meets this cone point, then it must be closed. In this case, we would have

$$\theta = \frac{\pi}{3(1 - 4\chi(\partial N))} = \frac{\pi}{4},$$

and we compute the lower bound for $\text{Vol}(N)$ as $0.428850\dots > H \cdot \text{Area}(\mathcal{T})$. Therefore, the shortest return path in N could not be contained in this order 4 axis. All the other cases follow similarly. In fact, we can rule out every possibility for a shortest return path in N *except* in the case that ∂N is a $(2, 4, 5)$ turnover and the shortest return path is contained in either the order 4 axis giving the order 4 cone point in ∂N or the order 5 axis giving the order 5 cone point in ∂N . We now rule these out by producing better lower bounds for $\text{Vol}(N)$ than those implied by Proposition 8.2.

In the order 4 case, observe that we now know exactly the radius of an embedded orbifold disk around the order 4 cone point, namely, the length of the segment from the order 4 cone point to the order 2 cone point in the $(2, 4, 5)$ turnover. This distance is given by $\cosh^{-1}(\sqrt{2} \cos(\pi/5))$. See Figure 6, which illustrates a lift of this disk to \mathbb{H}^2 . Referring to (8.4) in the proof of Proposition 8.2, we have that




FIGURE 6. A lift of the embedded orbifold disk of maximal radius around the order 4 cone point in a $(2, 4, 5)$ hyperbolic turnover

the length l of the shortest return path, if this path is contained in the order 4 axis, satisfies

$$\frac{\cosh l}{\cosh l - 1} \leq \cosh \left(2 \cosh^{-1} \left(\sqrt{2} \cos \frac{\pi}{5} \right) \right) = 4 \cos^2 \frac{\pi}{5} - 1.$$

This implies

$$\cosh l \geq \frac{4 \cos^2 \left(\frac{\pi}{5} \right) - 1}{4 \cos^2 \left(\frac{\pi}{5} \right) - 2}.$$

Since we know the relationship of the side length of a regular truncated 3-simplex to its angle (8.1), we can calculate θ so that l is at least the side length of T_θ . In this case, we have $\theta = 0.904556\dots$, and the corresponding lower bound for $\text{Vol}(N)$ is $0.383986\dots > H \cdot \text{Area}(\mathcal{T})$. So the shortest return path can not coincide with an order 4 cone point in ∂N .

We now consider the order 5 case. Fix a copy of $\pi_1(\mathcal{T})$ and let Π be the plane stabilized by it. By assumption, there is an elliptic element of order 5 which stabilizes a plane P and whose axis meets Π , and this axis projects to N as a shortest return path. Applying (6.2), we have that the axes of two elliptic isometries of order 5 must be at least $\delta(5, 5) \geq 0.736175\dots$ apart, or else they are equal or they intersect in a point, in order to generate a discrete group. In the last case, the two axes meet in a point which is stabilized by the group of isometries of a regular dodecahedron. Suppose that the axis A of an order 5 element intersects Π (and therefore corresponds to a local part of the immersion of \mathcal{T}), and that the cone angle associated to this element projects to the order 5 cone point in the $(2, 4, 5)$ turnover ∂N . Then there are several possibilities for the position of A . It cannot coincide with the axis of an order 4 element giving the order 4 cone point of

\mathcal{T} , for then this axis would have to correspond to an element of order at least 20, and by Table 1 this cannot occur. So A must be at least $\delta(4, 5) \geq 0.626869\dots$ away from any such order 4 axis. In this case, we must have that A is close enough to the axis of the order 5 elliptic giving the order 5 cone point of \mathcal{T} so that these axes either intersect or are equal.

Suppose first that they are equal, so that A gives the order 5 cone points in both \mathcal{T} and ∂N . Then the closed return path of shortest length, if it meets the order 5 cone point of ∂N , must connect these two cone points, one in \mathcal{T} and one in ∂N , and so its length is at least twice the length of this connecting segment (which is the unique common perpendicular for the plane P and the plane Π stabilized by our $(2, 4, 5)$ turnover subgroup), because the return path is closed. Now since the $(2, 4, 5)$ turnover ∂N is embedded, no other axis giving a cone point of \mathcal{T} can meet P , since two planes have a unique common perpendicular in \mathbb{H}^3 . This holds in particular for a closest order 4 axis to A . See Figure 7, which illustrates the plane containing these two axes. The points a and b correspond to order 5 and order 4 cone points




FIGURE 7. One way the shortest return path can meet the order 5 cone point in ∂N

of \mathcal{T} , respectively, and c to the order 5 cone point of ∂N . The segment between a and c is contained in the shortest return path, and the ray from b must be disjoint from the geodesic containing c . Then by the almost-right quadrilateral law of hyperbolic trigonometry, we have

$$L = \text{length}(\overline{ac}) > \sinh^{-1} \left(\frac{1}{\sinh(\text{length}(\overline{ab}))} \right).$$

Since \overline{ab} is the segment between cone points of order 4 and 5 in the $(2, 4, 5)$ turnover, its length is given by $\cosh^{-1}(\cot(\pi/5))$, so a calculation gives that L must be bigger than 0.921365.... Since the length of the shortest return path must be at least twice this, we can calculate the side length for the associated T_θ , where θ turns out to be 0.938037..., and the lower bound for $\text{Vol}(N)$ is $0.460222... > H \cdot \text{Area}(\mathcal{T})$. This proves that the shortest return path cannot connect the order 5 cone points of \mathcal{T} and ∂N .

The other possibility is that A intersects in a single point the order 5 axis giving the cone point of \mathcal{T} . In this case, this intersection point is stabilized by the orientation-preserving group of symmetries of the dodecahedron, and there is therefore an order 2 element whose axis meets A in an angle of $\cos^{-1}(\sqrt{2}/\sqrt{5 - \sqrt{5}})$. See Figure 8. But now we




FIGURE 8. The other way the shortest return path can meet the order 5 cone point in ∂N

can consider the common perpendicular segment between ∂N and this order 2 axis. The path along this common perpendicular which travels from ∂N to this order 2 and back again will give a shorter return path than the path contained in our order 5 axis. So the shortest return path could not have been contained in the order 5 axis.

This analysis implies that there can be no shortest return path in N , and so N could have had no boundary. We have therefore shown the following:

Proposition 9.1. *If a hyperbolic 3-orbifold Q contains an immersed, non-embedded $(2, 4, 5)$ turnover, then it can contain no embedded turnovers. In particular, $\text{Vol}(Q) < \text{Area}(\mathcal{T}(2, 4, 5)) = \pi/10 = 0.314159...$*

The last claim follows from the last statement in (5) of Theorem 1.1. The volume bound for such a hyperbolic 3-orbifold is rather

small, and there is such an orbifold [2]. Consider the tetrahedron from the Introduction, shown in Figure 1. Recall that Γ_3 denotes the orientation-preserving index two subgroup of the group generated by reflections in the faces of the this tetrahedron, and that the face ABC corresponds to an immersed $(2, 4, 5)$ turnover in the quotient orbifold $Q_3 = \mathbb{H}^3/\Gamma_3$. The volume of this orbifold is approximately 0.071770, which is less than a quarter of our volume bound. However, the $(2, 4, 5)$ turnover subgroup is actually contained in a $\mathbb{Z}/2\mathbb{Z}$ extension in $\pi_1(Q_3)$ (because it is contained in the group of reflections in the sides of the face ABC), and so, by the remarks following (7.1), the volume bound of the main theorem can actually be cut in half to $\pi/20 = 0.157079\dots$, which gives a reasonably close approximation for this example. The author conjectures that this is the only example of an immersed, non-embedded $(2, 4, 5)$ turnover in a hyperbolic 3-orbifold.

9.2. Immersed $(2, 4, 6)$ Turnovers. Consider the tetrahedron T_{10} in Figure 9. The notation for the dihedral angles is the same as in the previous example, and it can be similarly realized in \mathbb{H}^3 as the non-compact fundamental domain for the discrete group generated by reflections in its sides. It has one ideal vertex indicated at D . Let Γ_{10} be




FIGURE 9. A noncompact hyperbolic tetrahedron with $(2, 4, 6)$ triangular face

the orientation-preserving index two subgroup of the reflection group. This is also a Kleinian group, and we have $\text{Vol}(Q_{10} = \mathbb{H}^3/\Gamma_{10}) \approx 0.211446$. Now the face ABC of this tetrahedron is a $(2, 4, 6)$ triangle, and so Γ_{10} contains a $\mathbb{Z}/2\mathbb{Z}$ extension of a $(2, 4, 6)$ turnover group.

Clearly Q_{10} can contain no embedded turnovers, and so our volume bound is $\text{Area}(\mathcal{T}(2, 4, 6))/2 = 0.261799\dots$, which is very close to the volume of Q_{10} . We make the following conjecture:

Conjecture 9.2. *If a hyperbolic 3-orbifold Q contains an immersed, non-embedded $(2, 4, 6)$ turnover, then it can contain no embedded turnovers. In particular, $\text{Vol}(Q) < \text{Area}(\mathcal{T}(2, 4, 6)) = \pi/6 = 0.523598\dots$*

This conjecture is made in the place of calculating, as in the last example, that an immersed $(2, 4, 6)$ turnover can contain no embedded turnovers in its complement. The author believes that this can be shown, but the case by case analysis necessary is rather complicated.

9.3. Immersed $(2, 4, p)$ Turnovers for $p \geq 7$. The conjectural lack of embedded turnovers in the complement of an immersed $(2, 4, p)$ turnover does not extend to the case when $p \geq 7$. See Figure 10 (where $p = 764 \geq 7$ is arbitrary). The index two orientation-preserving sub-




FIGURE 10. A hyperbolic prism with immersed and embedded turnovers

group of the group generated by reflections in the sides of this prism is a Kleinian group by Andre'ev's Theorem (e.g. [28]). The corresponding orbifold contains an immersed $(2, 4, p)$ turnover (coming from the roof of the prism) and an embedded $(2, 3, p)$ turnover in the complement of this immersed turnover (coming from the base). (Recall from the introduction that Martin [21] showed that $(2, 3, p)$ turnovers are always embedded). Again, the fundamental group of the 3-orbifold $Q_{2,4,p}$ obtained from this prism contains a $\mathbb{Z}/2\mathbb{Z}$ extension of the $(2, 4, p)$ turnover group, and so the upper volume bound is

$\frac{H}{2}\text{Area}(\mathcal{T}(2, 4, p))$ ($H = 1.199678\dots$). This bound is best when $p = 7$, where we have

$$0.325947 \approx \text{Vol}(Q_{2,4,7}) < \frac{H}{2}\text{Area}(\mathcal{T}(2, 4, 7)) = 0.403810\dots,$$

and the bound worsens as p tends to infinity, where we have

$$0.501921 \approx \text{Vol}(Q_{2,4,\infty}) < \frac{H}{2}\text{Area}(\mathcal{T}(2, 4, \infty)) = 0.942225\dots$$

We observe that this family of examples demonstrates that the splitting collection of turnovers $\{\mathcal{T}_i\} \subset Q$ from Theorem 1.1 is not always empty. The fundamental groups of the orbifolds $Q_{2,4,p}$ in this section are examples of so-called *web* groups. Their limit sets form web-like Sierpinski gaskets on S_∞^2 with turnover groups for the component stabilizers.

9.4. Turnover Subgroups of Tetrahedral Groups. The principal reference for this subsection is Maclachlan [20], which classifies almost all the turnover subgroups contained in the cocompact hyperbolic tetrahedral groups. There are nine compact hyperbolic tetrahedra such that reflections in the sides of the tetrahedra tile hyperbolic 3-space [27, Chapter 7]. We denote them by $T_i[l_1, l_2, l_3; m_1, m_2, m_3]$, where i ranges from 1 to 9, and where the l_j and m_k give the integer submultiples of π at each edge AB, BC, AC, CD, AD, BD (in this order) for the tetrahedron $ABCD$ given in Figure 11. Let Γ_i denote




FIGURE 11. The tetrahedron $ABCD$

the orientation-preserving index two subgroup of the group generated by the reflections in the faces of T_i . Then the Γ_i are Kleinian groups which are all arithmetic except for Γ_8 , and Maclachlan classifies most

of the turnover subgroups for the eight arithmetic examples by arithmetic means. We will see that it is possible, in principle, to locate these turnover subgroups by purely geometric means. We will focus on $T_8[2, 3, 4; 2, 3, 5]$ and $T_9[2, 3, 5; 2, 3, 5]$. If we have a turnover subgroup $T(p, q, r) < \Gamma_i$, then the plane Π_T invariant under T will always intersect at least three intersecting elliptic axes in \mathbb{H}^3 (corresponding to rotations in three edges of the tetrahedron T_i). Therefore, at least one of these axes must meet Π_T non-perpendicularly. It follows that any turnover subgroup $T(p, q, r) < \Gamma_i$ will yield an immersed turnover in the quotient 3-orbifold.

We begin by correcting an error in Maclachlan's work. He claims to prove [20, Theorem 5.1] that $\mathcal{O}_9 = \mathbb{H}^3/\Gamma_9$ contains the turnovers $(3, 3, 5)$ and $(5, 5, 5)$. He also states that \mathcal{O}_9 *may* contain a $(3, 5, 5)$ turnover. The volume of \mathcal{O}_9 can be computed as approximately 1.004261. Since \mathcal{O}_9 contains no embedded turnovers, the upper bound for its volume given by Theorem 1.1 is just the area of any immersed turnover that it contains. But the volume bound in the case of the $(3, 3, 5)$ turnover is $4\pi/15 = 0.837758... < 1.004261$. We conclude that \mathcal{O}_9 can contain no $(3, 3, 5)$ turnover. Maclachlan's claim that \mathcal{O}_9 contains a $(5, 5, 5)$ turnover is based on the fact that $T(5, 5, 5) < T(3, 3, 5)$, and so this claim becomes unjustified.

As it turns out, we can prove directly that \mathcal{O}_9 contains both a $(3, 5, 5)$ turnover and a $(5, 5, 5)$ turnover. We begin with the $(3, 5, 5)$ turnover. See Figure 12, which illustrates one fundamental domain for Γ_9 in \mathbb{H}^3 . The numbers at the edges represent the submultiples




FIGURE 12. A fundamental domain for Γ_9 in \mathbb{H}^3

of π of the dihedral angle at that edge (so, for instance, the edge BD is labeled $5/2$ because the dihedral angle at that edge is $2\pi/5$).

We remark that this illustration is, obviously, not in true perspective. Note that each edge also represents the axis of an elliptic isometry in Γ_9 , whose order is the numerator of the edge label. We will prove the claim by finding three planes whose pairwise lines of intersection are composed of axes of elliptic elements that generate a $(3, 5, 5)$ turnover subgroup.

Let $\angle(X, Y)$ denote the dihedral angle of the intersecting planes X and Y . It is easily seen from the figure that $\angle(ACD, ACC') = \pi/5$ and $\angle(ACC', BC'D) = \pi/3$, and that these pairs of planes intersect in order 5 and order 3 elliptic axes, respectively. The planes $BC'D$ and ACD contain D in their intersection, and we can calculate their line of intersection and the dihedral angle as follows. The point D is stabilized by a $(2, 3, 5)$ spherical turnover subgroup $G \leq \Gamma_9$. The plane ACD is spanned by lines $L_1 \supset CD$ and $L_2 \supset AD$, which are the axes of elements in G of order 2 and 3, respectively, and which are as close as possible in terms of their angle of intersection. Similarly, from the plane $BC'D$ we obtain closest possible axes $L_3 \supset C'D$ and $L_4 \supset BD$ of elements of order 2 and 5, respectively. To calculate the angle $\angle(ACD, BC'D)$, then, we simply calculate the angle between the planes spanned by these axes in the spherical link of D . See Figure 13. The image at left is just Figure 12 with the axes L_i labeled,




FIGURE 13. The intersection of planes in the link of the vertex D

and the image at right shows two faces of the spherical dodecahedron around D , corresponding to the $(2, 3, 5)$ spherical turnover subgroup G . The labels L_i are meant to indicate the intersection of L_i with the link. Now it is easily seen that $\angle(ACD, BC'D) = \pi/5$, and that this intersection occurs along the axis of an order 5 elliptic element. So we have three planes which pairwise intersect in angles giving a $(3, 5, 5)$

The same method can also be used to locate a $(5, 5, 5)$ turnover in \mathcal{O}_9 . See Figure 14, which illustrates two copies (that is, four tetrahedra) of a fundamental domain for the action of Γ_9 . Since the edge




FIGURE 14. Two fundamental domains for Γ_9 in \mathbb{H}^3

labels indicate the submultiples of π for the polyhedron pictured, we observe from the figure that the points A, B, C, C' , and D' are coplanar, and similarly for A', B, D , and D' . It is then easily seen from the figure that we have $\angle(ACD', ACD) = \angle(ACD', A'DD') = \pi/5$. Additionally, by the method illustrated in Figure 13, it is readily seen that $\angle(ACD, A'DD') = \pi/5$. All of these intersections occur along order 5 axes of elements in Γ_9 , and so as in the last case we conclude that \mathcal{O}_9 contains a $(5, 5, 5)$ turnover.

The method just described can be used to verify many of the examples of turnover subgroups in tetrahedral groups given by Maclachlan. It is also possible, in principle, to describe a geometric algorithm to locate all the examples. Because the lengths of the sides of the possible turnovers in these tetrahedral groups are bounded, and because the dihedral edge at each elliptic axis in these tetrahedra is fixed by

a translation in the universal cover \mathbb{H}^3 , the search for axes which are mutually at the appropriate distance apart from one another will be a search over a compact subset of hyperbolic 3-space. We will not construct an explicit algorithm. However, we note that the only possibilities for turnover subgroups of Γ_9 are $(2, 5, 5)$, $(3, 3, 5)$, $(3, 5, 5)$, and $(5, 5, 5)$ (because of the orders of the edges of T_9), and since we can rule out the first two of these (the $(2, 5, 5)$ case follows by our volume bound, exactly as in the $(3, 3, 5)$ case) and prove the existence of the other two, we have the following:

Proposition 9.3. *The turnover subgroups of the group Γ_9 are $T(3, 5, 5)$ and $T(5, 5, 5)$.*

Turning to T_8 , we have $\text{Vol}(\mathcal{O}_8 = \mathbb{H}^3/\Gamma_8) \approx 0.717306$. The orders of the cone points for any turnover in \mathcal{O}_8 must come from the dihedral angles at the edges of T_8 , and so the list of possible turnovers is

$$(2, 4, 5), (2, 5, 5), (3, 3, 4), (3, 3, 5), (3, 4, 4),$$

$$(3, 4, 5), (3, 5, 5), (4, 4, 4), (4, 4, 5), (4, 5, 5), (5, 5, 5).$$

The volume bound of the main theorem rules out the left-most three turnovers from the top row. Of those that remain, only $(3, 4, 5)$ is non-arithmetic. It can be shown that \mathcal{O}_8 contains this turnover, using the exact same analysis (and the same figures with slightly different edge labels) as in the $T(3, 5, 5) < \Gamma_9$ case above.

As in the $T(5, 5, 5) < \Gamma_9$ case, one can prove the existence of a $(4, 5, 5)$ turnover in \mathcal{O}_8 as well. Once again, the proof is identical, and the figure required to carry out the proof differs from Figure 14 only in some of the edge labels.

The group Γ_8 may contain other arithmetic turnover subgroups from the list above. However, in the absence of an explicit algorithm (whose existence was described above), we note that the presence of the arithmetic turnover subgroup $T(4, 5, 5)$ provides interesting information about the orbifold \mathcal{O}_8 , due to a recent result of Long-Lubotzky-Reid [19, Proposition 4.1]:

Proposition 9.4. *The orbifold \mathcal{O}_8 cannot contain immersed turnovers of type $(2, 4, 5)$, $(2, 5, 5)$, or $(3, 3, 4)$. It contains immersed $(3, 4, 5)$ and $(4, 5, 5)$ turnovers, and the existence of the latter implies that \mathcal{O}_8 has a tower of principle congruence covers with Property τ .*

10. CALCULATIONS

In this section, we will prove some results about hyperbolic turnovers and thrice-punctured spheres, as well as determine an interesting new isoperimetric constant. Our first result is a bound on the diameter of a maximally embedded disk in a turnover.

Proposition 10.1. *Let $\mathcal{T}(p, q, r)$ be a hyperbolic turnover. Then the radius of an embedded disk in \mathcal{T} is less than*

$$(10.1) \quad \ln \left(\frac{2 + \sqrt{7}}{\sqrt{3}} \right).$$

Proof of 10.1. We will prove the result by proving that the maximal radius r_{max} of an embedded disk in a hyperbolic thrice-punctured sphere is given by (10.1), and that the maximal radius of an embedded disk in a hyperbolic turnover must be less than r_{max} .

We begin with the thrice-punctured sphere S . Consider a fundamental domain in \mathbb{H}^2 for S , normalized as in Figure 15. We have para-




FIGURE 15. A fundamental domain for a thrice-punctured sphere, with three isometric horospheres

bolic isometries which fix 0, -2 , and ∞ , the latter acting on the upper half-space as translation by 4. Denote this translation by $g: z \mapsto z+4$. It is an easy computation to show that a rotation σ thru $2\pi/3$ radians centered at the “orthocenter” $(-1, \sqrt{3})$ of the ideal triangle on the left permutes the points 0, -2 , ∞ cyclically. We consider the horoballs based at 0, -2 , and ∞ , each passing through $(-1, \sqrt{3})$, as in the figure. These horoball regions are permuted by σ with their centers, and any point in S lifts to a point contained in at least one of these

horoballs. It is therefore the case that moving in any direction away from $(-1, \sqrt{3})$ is equivalent, up to the action of σ , to increasing the Euclidean height from $\sqrt{3}$. However, if a point in the upper half-space has coordinates (a, b) , then a computation shows that the distance to its translate by g is given by

$$2 \ln \left(\frac{2 + \sqrt{b^2 + 4}}{b} \right),$$

which is less than twice the quantity in (10.1) when $b > \sqrt{3}$. Any point in the fundamental domain for S is taken to its closest translate by the action of g , up to the isometry σ . So the largest embedded disk in S lifts to a disk based at $(-1, \sqrt{3})$. But, as observed above, a disk based at $(-1, \sqrt{3})$ will meet its closest translates by $\pi_1(S)$ *exactly* when it has radius given by (10.1).

To complete the proof, we construct a contraction from the thrice-punctured sphere to the smooth part of a hyperbolic turnover. Consider a triangle Δ in the projective model of \mathbb{H}^2 , positioned so that its circumcenter lies at the center of the disk. Then the vertices of Δ lie on a circle of that is concentric with the unit disk, and Δ is the (Euclidean) homothetic image of an ideal triangle. It is easily seen that the Euclidean homothety taking this ideal triangle to Δ is a contraction for the hyperbolic metric. If we double along the boundary of Δ , then this gives a contraction from the thrice-punctured sphere to the double of Δ that restricts to a homeomorphism on the smooth part of the double of Δ .

Now if the maximal radius of an embedded ball in a hyperbolic turnover is greater than or equal to r_{max} , then we have an immediate contradiction. For consider the center p of such a ball, and let ℓ be the shortest loop through p . Then the length of ℓ is at least $2r_{max}$. But then pulling ℓ back to the thrice-punctured sphere, decreasing the length of the resulting loop (which we can do because the thrice-punctured sphere has maximal injectivity radius r_{max}), and contracting back to the turnover will produce a shorter path through p . This proves the proposition. 10.1

Our final calculation provides the volume bound of a nice room in terms of the area of its ceiling, which is used in Corollary 4.2. In the terminology of Section 4, let F be a floor in a hyperbolic plane $\Pi \subset \mathbb{H}^3$, and let B and S be a nice room over F and its ceiling, respectively.

Let H be the height of B . Recall that the metric for \mathbb{H}^3 using Fermi coordinates based on Π is given by $dh^2 + \cosh^2 h(dr^2 + \sinh^2 r d\theta^2)$, where r and θ describe polar coordinates on Π and h is perpendicular distance to Π . Write $dA = \sinh^2 r dr d\theta$ for the area form on Π . We have

$$(10.2) \quad \frac{\text{Area}(S)}{\text{Vol}(B)} = \frac{\iint_F \cosh^2 H dA}{\iint_F \int_0^H \cosh^2 h dh dA} = \frac{4 \cosh^2 H}{\sinh 2H + 2H}.$$

The function of H given above has interesting properties. It limits to infinity as H approaches zero, and to 2 as H approaches infinity. It is *not* monotonic, however. It has a minimum value of $2/H = 1.667113\dots$ when H is the positive solution of $\coth x = x$, and it decreases from infinity to this value and then increases to 2 from this value. As a side remark, the author would be very interested in understanding the significance of the critical point of this function, as it seems to defy intuition for this ratio to be non-monotone. Figure 16 shows a graph of this function. In particular, we have

$$\frac{\text{Area}(S)}{\text{Vol}(B)} \geq 1.667113\dots,$$

which is the result necessary for Corollary 4.2.

REFERENCES

- [1] I. Agol, *Tameness of hyperbolic 3-manifolds*, 2004, [arXiv:math.GT/0405568](#).
- [2] T. Baskan and A. M. Macbeath, *Centralizers of reflections in crystallographic groups*, Math. Proc. Cambridge Philos. Soc. **92** (1982), no. 1, 79–91. MR 662962
- [3] Michel Boileau, Bernhard Leeb, and Joan Porti, *Geometrization of 3-dimensional orbifolds*, Ann. of Math. (2) **162** (2005), no. 1, 195–290. MR 2178962
- [4] Michel Boileau, Sylvain Maillot, and Joan Porti, *Three-dimensional orbifolds and their geometric structures*, Panoramas et Synthèses [Panoramas and Syntheses], vol. 15, Société Mathématique de France, Paris, 2003. MR 2060653
- [5] K. Böröczky, *Packing of spheres in spaces of constant curvature*, Acta Math. Acad. Sci. Hungar. **32** (1978), no. 3-4, 243–261. MR 512399
- [6] Yu. D. Burago and V. A. Zalgaller, *Geometric inequalities*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 285, Springer-Verlag, Berlin, 1988, Translated from the Russian by A. B. Sosinskiĭ, Springer Series in Soviet Mathematics. MR 936419
- [7] Danny Calegari, *Napoleon in isolation*, Proc. Amer. Math. Soc. **129** (2001), no. 10, 3109–3119 (electronic). MR 1840118
- [8] Isaac Chavel, *Riemannian geometry: A modern introduction*, second ed., Cambridge Studies in Advanced Mathematics, vol. 98, Cambridge University Press, Cambridge, 2006. MR 2229062





FIGURE 16. The function $H \mapsto \frac{4 \cosh^2 H}{\sinh 2H + 2H}$


- [9] Daryl Cooper, *The volume of a closed hyperbolic 3-manifold is bounded by π times the length of any presentation of its fundamental group*, Proc. Amer. Math. Soc. **127** (1999), no. 3, 941–942. MR 1646313
- [10] Daryl Cooper, Craig D. Hodgson, and Steven P. Kerckhoff, *Three-dimensional orbifolds and cone-manifolds*, MSJ Memoirs, vol. 5, Mathematical Society of Japan, Tokyo, 2000, With a postface by Sadayoshi Kojima. MR 1778789
- [11] William D. Dunbar, *Hierarchies for 3-orbifolds*, Topology Appl. **29** (1988), no. 3, 267–283. MR 953958
- [12] William D. Dunbar and G. Robert Meyerhoff, *Volumes of hyperbolic 3-orbifolds*, Indiana Univ. Math. J. **43** (1994), no. 2, 611–637. MR 1291531
- [13] D. B. A. Epstein and A. Marden, *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces*, Analytical and geometric aspects of hyperbolic space (Conventry/Durham, 1984), London Math. Soc. Lecture Note Ser., vol. 111, Cambridge Univ. Press, Cambridge, 1987, pp. 113–253. MR 903852
- [14] F. W. Gehring and G. J. Martin, *Commutators, collars and the geometry of Möbius groups*, J. Anal. Math. **63** (1994), 175–219. MR 1269219
- [15] ———, *On the minimal volume hyperbolic 3-orbifold*, Math. Res. Lett. **1** (1994), no. 1, 107–114. MR 1258496
- [16] ———, *(p, q, r) -Kleinian groups and the Margulis constant*, Complex analysis and dynamical systems II, Contemp. Math., vol. 382, Amer. Math. Soc., Providence, RI, 2005, pp. 149–169. MR 2175884
- [17] R. Kirby, *Problems in low-dimensional topology*, 1996.
- [18] Sadayoshi Kojima, *Deformations of hyperbolic 3-cone-manifolds*, J. Differential Geom. **49** (1998), no. 3, 469–516. MR 1669649
- [19] D. D. Long, A. Lubotzky, and A. W. Reid, *Heegaard genus and Property τ for hyperbolic 3-manifolds*, Journal of Topology **1** (2008), 152–158.
- [20] C. Maclachlan, *Triangle subgroups of hyperbolic tetrahedral groups*, Pacific J. Math. **176** (1996), no. 1, 195–203. MR 1433988
- [21] Gaven J. Martin, *Triangle subgroups of Kleinian groups*, Comment. Math. Helv. **71** (1996), no. 3, 339–361. MR 1418942


- [22] Bernard Maskit, *Kleinian groups*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 287, Springer-Verlag, Berlin, 1988. MR 959135
- [23] William H. Meeks, III and Shing Tung Yau, *Topology of three-dimensional manifolds and the embedding problems in minimal surface theory*, Ann. of Math. (2) **112** (1980), no. 3, 441–484. MR 595203
- [24] Robert Meyerhoff, *A lower bound for the volume of hyperbolic 3-manifolds*, Canad. J. Math. **39** (1987), no. 5, 1038–1056. MR 918586
- [25] Yosuke Miyamoto, *Volumes of hyperbolic manifolds with geodesic boundary*, Topology **33** (1994), no. 4, 613–629. MR 1293303
- [26] Walter D. Neumann and Alan W. Reid, *Rigidity of cusps in deformations of hyperbolic 3-orbifolds*, Math. Ann. **295** (1993), no. 2, 223–237. MR 1202390
- [27] John G. Ratcliffe, *Foundations of hyperbolic manifolds*, Graduate Texts in Mathematics, vol. 149, Springer-Verlag, New York, 1994. MR 1299730
- [28] Roland K. W. Roeder, John H. Hubbard, and William D. Dunbar, *Andreev's theorem on hyperbolic polyhedra*, Ann. Inst. Fourier (Grenoble) **57** (2007), no. 3, 825–882. MR 2336832
- [29] David Singerman, *Finitely maximal Fuchsian groups*, J. London Math. Soc. (2) **6** (1972), 29–38. MR 0322165
- [30] W. P. Thurston, *The geometry and topology of 3-manifolds*, Lecture notes from Princeton University, 1978–80.


DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, FAIRFIELD UNIVERSITY, FAIRFIELD, CT 06824

E-mail address: srafalski@fairfield.edu










(a)



(b)

